

60
00/4T
U.S.
PTO

Express Mail Label No. EL680251845US

9-15-00

A

UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
13862(YOR920000575US1)

Total Pages in this Submission
3

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

A MECHANISM FOR SYNCHRONOUS INTERPROCESS COMMUNICATION OVER TRANSPARENT EXTERNAL MONITORS

and invented by:

Trent Ray Jaeger
Jonathon Earnshaw Tidswell

Jc918 U.S. PTO
09/661341 PTO
09/14/00

If a CONTINUATION APPLICATION, check appropriate box and supply the requisite information:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. Filing fee as calculated and transmitted as described below
2. Specification having 25 pages and including the following:
 - a. Descriptive Title of the Invention
 - b. Cross References to Related Applications (*if applicable*)
 - c. Statement Regarding Federally-sponsored Research/Development (*if applicable*)
 - d. Reference to Microfiche Appendix (*if applicable*)
 - e. Background of the Invention
 - f. Brief Summary of the Invention
 - g. Brief Description of the Drawings (*if drawings filed*)
 - h. Detailed Description
 - i. Claim(s) as Classified Below
 - j. Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
13862(YOR920000575US1)

Total Pages in this Submission
3

Application Elements (Continued)

3. Drawing(s) (*when necessary as prescribed by 35 USC 113*)
a. Formal Number of Sheets 15
b. Informal Number of Sheets _____
4. Oath or Declaration
a. Newly executed (*original or copy*) Unexecuted
b. Copy from a prior application (37 CFR 1.63(d)) (*for continuation/divisional application only*)
c. With Power of Attorney Without Power of Attorney
d. DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (*usable if Box 4b is checked*)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied
under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.
6. Computer Program in Microfiche (*Appendix*)
7. Nucleotide and/or Amino Acid Sequence Submission (*if applicable, all must be included*)
a. Paper Copy
b. Computer Readable Copy (*identical to computer copy*)
c. Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. Assignment Papers (*cover sheet & document(s)*)
9. 37 CFR 3.73(B) Statement (*when there is an assignee*)
10. English Translation Document (*if applicable*)
11. Information Disclosure Statement/PTO-1449 Copies of IDS Citations
12. Preliminary Amendment
13. Acknowledgment postcard
14. Certificate of Mailing

First Class Express Mail (*Specify Label No.*): EL680251845US

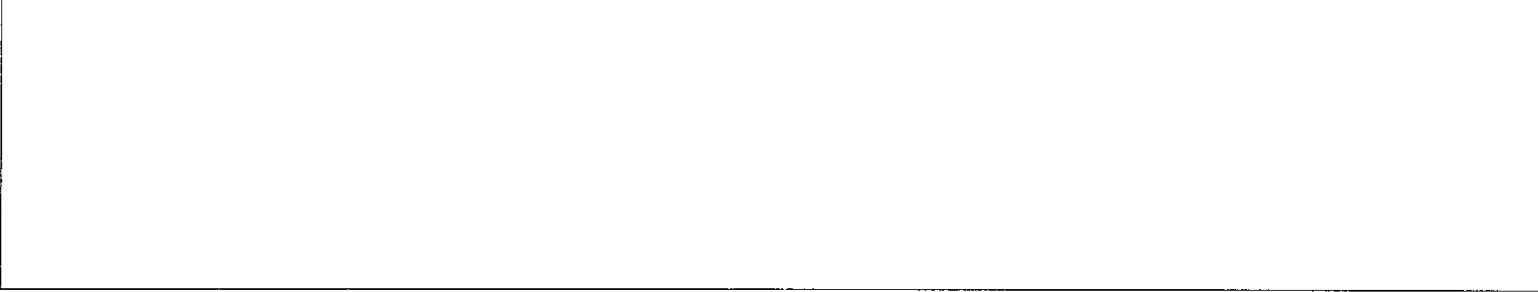
UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
13862(YOR920000575US1)

Total Pages in this Submission
3

Accompanying Application Parts (Continued)

15. Certified Copy of Priority Document(s) (*if foreign priority is claimed*)
16. Additional Enclosures (*please identify below*):


Fee Calculation and Transmittal

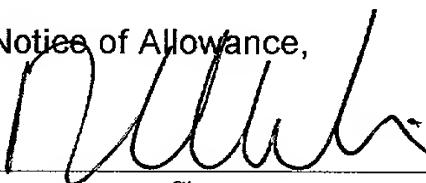
CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	35	- 20 =	15	x \$18.00	\$270.00
Indep. Claims	4	- 3 =	1	x \$78.00	\$78.00
Multiple Dependent Claims (check if applicable)	<input type="checkbox"/>				\$0.00
				BASIC FEE	\$690.00
OTHER FEE (specify purpose)					\$0.00
				TOTAL FILING FEE	\$1,038.00

A check in the amount of _____ to cover the filing fee is enclosed.

The Commissioner is hereby authorized to charge and credit Deposit Account No. **50-0510/IBM** as described below. A duplicate copy of this sheet is enclosed.

- Charge the amount of **\$1,038.00** as filing fee.
- Credit any overpayment.
- Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).



Signature

Richard L. Catania
Registration No. 32,608

SCULLY, SCOTT, MURPHY & PRESSER
400 Garden City Plaza
Garden City, New York 11530
(516)742-4343

Dated: September 14, 2000

cc:

A MECHANISM FOR SYNCHRONOUS INTERPROCESS
COMMUNICATION OVER TRANSPARENT EXTERNAL MONITORS

Background Of The Invention

5

Field of Invention

10

The present invention relates generally to the field of interprocess communication. More particularly, this invention relates to an interprocess communication mechanism that enables the synchrony semantics of each communication to be controlled by external entities called monitors, such that: (1) communication may be redirected to such external entities transparently to the sender and receiver of the communication and (2) the monitors can determine the semantics of a completed communication as they desire.

15
20
25
30

Background of the Invention

25

Interprocess communication (IPC) monitoring enables the examination of any IPC between a source and a destination. IPC monitoring is useful for a variety of purposes, including debugging, logging, and security. For example, a monitor may collect communication state for the purpose of debugging a program consisting of several independent tasks. Also, a monitor can be used to filter communication data or control the communication rate for security purposes.

30

Transparent monitoring means that the source and destination are not aware that they are being monitored.

YOR920000575US1

This has two advantages: (1) the system can control the insertion and removal of monitors without interacting with either the source or destination, and (2) the source and destination protocols do not need to take into account the possibility that they may be monitored.

Traditional systems make no attempt to support transparent monitoring. Using Mach-style ports (see K.

Loepere, Mach 3 Kernel Principles, Open Software

Foundation and Carnegie Mellon University, 1992), the

source and destination hold rights that must be revoked in order to insert a monitor, so at least one must be notified before such a change can occur safely. The

Pebble microkernel enables transparent redirection of source IPCs using its portals to implement customized IPC, but the redirection is not transparent to the destination because it sees that the message is from the redirected task, not the original

(see E. Gabber et al, Building Efficient Operating Systems from User-level Components in Pebble, in Proceedings of the 1999 USENIX Annual Technical conference, 1999).

Other IPC mechanisms, such as Clans & Chiefs (see J. Liedtke, Clans & Chiefs, in Architektur von Rechensystemen, 1992, in English) and IPC Redirection (see, "Flexible Access Control Using IPC Redirection," HotOS 1999), enable monitors to intercept and forward IPCs while claiming to be the original source of the IPC. Thus, the destination receives the IPC from the source, not the monitor, so it need not know that an IPC is being monitored. Unfortunately, such mechanisms are not truly transparent because the kernel's IPC semantics are not

5

10

15
20
25

30

preserved when a monitor intercepts an IPC. Modern microkernels implement a synchronous IPC semantics, which means that the source is blocked until the destination is ready to receive the IPC or an error occurs (e.g., the destination task is killed or a timeout expires). When destination commences receipt, the IPC is sent to the destination and the source unblocks. Unfortunately, if a monitor is inserted on an IPC path, the source is unblocked when the IPC is received by the monitor, not the destination. This may result in some anomalous behaviors, such as: (1) the source assuming that IPCs have been delivered to the destination before they really have; (2) the source terminating IPCs due to timeout expiration even though the destination is ready, but because the monitor was not ready; and (3) the source assuming that the IPC was delivered reliably to the destination when an error may have occurred.

What is needed is that the system (i.e., the kernel and the monitors) control the synchrony of each communication, so that the communication appears to the sender and destination to be implemented according to the same semantics regardless of whether a monitor is present or not. Further, the system may choose arbitrary actions upon a communication, so the interprocess communication mechanism must permit the system to provide any desired semantics. Lastly, the interprocess communication mechanism must result in a system that is robust in the presence of errors.

Summary of the Invention

An object of this invention is to improve interprocess communications.

5

Another object of the present invention is to provide a method and system for controlling the synchrony of interprocess communications so that the communications appear to the senders and destinations to be implemented according to the same semantics regardless of whether the communications are monitored or not.

10

A further object of this invention is to use monitors, for monitoring interprocess communications, as extensions of a system kernel, to that, when a particular communication is monitored, the source and destination of the communication are treated as if the kernel is still processing the communication.

15
20
25
30

These and other objects are attained with a method and system for performing interprocess communications (IPCs). The method comprises the steps of receiving IPC requests, where each of the IPC requests identifies a source and a destination; building IPCs in response to the request; transmitting the IPCs from the sources to the destinations; and intercepting and examining selected ones of the IPCs. The method comprises the further step of controlling the synchrony of the IPCs so that each IPC appears to its source and destination to be implemented according to the same semantics regardless of whether the IPC is intercepted and examined.

With the preferred embodiment of this invention, the system monitors are considered as an extension of the system kernel (although they may be linked into the kernel and run in kernel mode as well), so the source and destination are treated as if the kernel is still processing the IPC. Thus, the desired semantics of communication can be implemented in the monitors. However, there are a number of possible monitoring semantics that may be reasonable, so in order to maintain utility of the approach, the IPC mechanism enables the monitors to coordinate using a small set of actions.

These actions are as follows. First, the interception of an IPC by a monitor permits the monitor to manage: (1) the identity of the source that ultimately must be unblocked; (2) when the sender is unblocked; and (3) whether the monitor is notified when the communication reaches the destination. In addition, error handling semantics can be expressed by the monitors such that all communication state in the monitors is maintained consistently.

Further benefits and advantages of the invention will become apparent from a consideration of the following detailed description, given with reference to the accompanying drawings, which specify and show preferred embodiments of the invention.

Brief Description of the Drawings

Figure 1 displays the general mechanism for delivering IPCs synchronously regardless of whether one or more

monitors are inserted on the IPC path between the source and destination.

5 Figure 2 displays how the delivery mechanism for IPCs is chosen.

Figure 3 displays how the kernel implements a source to monitor IPC delivery.

10 Figure 4 displays how the kernel implements a monitor to monitor IPC delivery.

Figure 5 displays how the kernel implements a monitor to destination IPC delivery.

15 Figure 6 displays the process by which a monitor handles IPCs that are redirected to it while managing the synchronicity information.

20 Figure 7 displays how the monitor prepares an IPC that it is forwarding.

Figure 8 displays how an IPC error message is processed.

25 Figure 9 displays how an IPC unblock message is processed.

Figure 10 displays how the synchrony is implemented when the IPC is delivered to the destination.

30 Figure 11 displays how an IPC error message is built.

Figure 12 displays how an IPC error message is delivered.

Figure 13 displays how an IPC unblock message is built.

5 Figure 14 displays how an IPC unblock message is
delivered.

Figure 15 displays how an IPC notify message is built.

10 Figure 16 displays how an IPC notify message is
delivered.

15 Figure 17 displays how the destination to which a source
is blocked is changed should a monitor redirect the
message to an alternative final destination.

Detailed Description of the Preferred Embodiment

20 For the preferred embodiment, the L4 microkernel's IPC
mechanism as the base IPC mechanism. The L4 IPC
mechanism is described in the L4 Reference Manual. L4
enables the creation of tasks that can be inserted on
arbitrary IPC paths. A mechanism called IPC redirection
is used in the preferred embodiment to determine whether
an IPC is sent to a monitor or to a destination
25 ("Flexible Access Control Using IPC Redirection," HotOS
1999). One suitable redirection mechanism is disclosed
in copending patent application no. 09/276,869, filed
March 26, 1999, entitled "Flexible Interprocess
30 Communication via Redirection," the disclosure of which
is herein incorporated by reference.

Figure 1 shows the basic synchronous IPC mechanism that enables transparent use of an arbitrary number of monitors. First, the source requests an IPC send (100). The sender (in this case, the source) is then blocked (110). The kernel processes the IPC delivery (200) which results in the IPC being sent either to the destination or a monitor (120). The kernel uses the IPC redirection mechanism to determine whether the IPC is delivered to the destination or a monitor. If the IPC is delivered to the final destination, then the kernel completes the synchronous IPC processing to ensure that the source can be unblocked (1000). If the IPC is delivered to a monitor, then the monitor performs its processing and maintains the synchrony information of the IPC (700). The monitor then decides whether to forward the IPC to the destination (or an alternative destination) or not (130). If the IPC is forwarded, the monitor sends the IPC using the mechanisms described in (200). Otherwise, the monitor decides whether to return an error (140). Error messages are a new type of IPC, so a new mechanism is described for handling them (800). If no error is returned, the monitor simply waits for a subsequent request (150).

Figure 2 describes how the IPC delivery mechanism is chosen (200). This mechanism is executed by the kernel in the preferred embodiment. First, the IPC redirection mechanism determines the immediate destination of an IPC (210) (220): (1) an invalid destination indicates that the IPC is to be blocked (230); (2) the IPC may be sent directly to a destination; or (3) the IPC may be redirected to a monitor. In addition to the L4 IPC

5

10

redirection mechanism, monitor tasks must be signified to the kernel, so completion processing can proceed (1000). In either case, the IPC may be deceived (i.e., sent from a task claiming to be another (240), and the delivery mechanism choice depends on whether a monitor or the source directly sent the IPC (250). Different mechanisms are applied for: (1) source to destination delivery (260) (this is the standard L4 IPC delivery mechanism, so it is not covered); (2) source to monitor delivery (300); monitor to monitor delivery (400); or monitor to destination delivery (500).

15
20

25

Figure 3 describes how the kernel processes an IPC delivery from a source to a monitor (300). The kernel simply delivers the IPC from the source to the monitor (310) and unblocks the monitor (320). This delivery (310) is performed using the traditional deceiting IPC as described in the L4 Users' Manual. After the IPC is delivered, only the monitor is unblocked (320). The source remains blocked.

30

Figure 4 describes how the kernel processes an IPC delivery from a monitor to another monitor (400). First, the kernel sets the traditional deceiting IPC data for the IPC as described in the L4 Users' Manual (410). Then, the kernel determines the type of the IPC (420) and sets the IPC type value appropriately (430). The various IPC type values are listed below:

IPC type values

direct IPC:	0
error:	1
notification:	2
unblocking:	3
deceit:	4
deceit w/ blocked source:	5
deceit w/ controlling monitor:	6
Both 5 and 6:	7

The kernel then pushes any synchrony values onto the stack of the IPC receiver (440). If the IPC type is 5, then only the blocked source is pushed. If the IPC type is either 6 or 7, then both the blocked source and controlling monitor values are pushed. If the IPC type is 0-4 no parameters are pushed. The kernel then delivers the IPC using the traditional L4 mechanism (450). Both the sender and receiver of the IPC are unblocked (460) (470) although according to L4 IPC semantics, the receiver is given the remainder of the sender's timeslice, so it become the next runnable task (480).

Figure 5 describes how the kernel sends an IPC to the destination (500). Basically, the mechanism is the same except that there is no need for the synchrony information to be passed to the destination. Therefore, IPC types 5-7 are equivalent to 4 in this case. First, the IPC is prepared for a deceived send as before (410). Then, the IPC type given that the IPC is to be delivered to the destination is determined (510). Then, the IPC

type is set (430). Next, the IPC is delivered by the kernel as is traditional in L4 (450). Then, the sender and receiver are unblocked (460) (470). Lastly, the destination is made runnable (480).

5
Figure 6 describes how the monitor processes the IPCs that are redirected to it (700). IPCs that are intended for the monitor are processed as traditional L4 IPCs.
10 First, the monitor receives the IPC message and error codes (710). Using this information, the monitor can do arbitrary, monitor-specific processing (720). If the monitor detects an error (720) (either through the error codes or its monitor-specific processing), the monitor then processes the IPC error (800). If there is no error, then the monitor determines whether to unblock the source (740). If the source is to be unblocked (750), the monitor must process a source unblock (900).
15 Regardless, the monitor creates a deceiting IPC message typical to L4 except that the monitor can now also specify the blocked source and/or the controlling monitor (760). Lastly, the monitor determines whether to update the IPC block state of the original source (770).

20
25 Figure 7 describes how the monitors prepare IPCs that they are going to forward. First, the monitor sets the blocked source (762) and controlling monitor (764) values for the IPC as necessary. The blocked source need only be set (762) if the identity of the original source of the IPC is different than the true source specified in the deceived IPC. The monitor may set the value of the controlling monitor to (764): (1) null, if none is specified; (2) itself if it wants to receive notification
30

5

when the IPC has been completed (i.e., been delivered to the destination); or (3) the current controlling monitor value it received in the incoming IPC. Subsequently, the monitor completes and delivers a typical L4 deceived IPC (766) and then waits for further redirected IPCs (150).

10

Figure 8 describes how an IPC error message is processed (800). First, the monitor builds an IPC error message (810). Second, the monitor asks the kernel to deliver the IPC error message using the basic L4 IPC mechanism (820). The IPC error message is itself an IPC, so it can be delivered using a slight generalization of the basic IPC mechanism. Then, the monitor waits for the next IPC (150).

10

Figure 9 describes how an IPC unblock message is processed (900). First, the monitor builds an IPC unblock message (910). Next, the monitor asks the kernel to deliver the IPC unblock message (920). Like the IPC error message case, the IPC unblock message is also an IPC. Lastly, the monitor waits for the next IPC (150).

25

30

Figure 10 describes the additional IPC effort when the IPC is delivered to the destination from the monitor (1000). At this time, the kernel must determine how to proceed with unblocking the source. If a controlling monitor is set (1010) (1020), then a notify message is constructed which is used to inform the controlling monitor that IPC from the source has been completed (1030). This message is delivered to the controlling monitor (1040). If no controlling monitor is set, the kernel then determines whether a blocked source is set

(1050) (1060). If so, the blocked source task is unblocked (1070). Otherwise, the value of the true source is unblocked (1070).

5 Figure 11 describes the mechanism for building error messages (810). The key component is the setting of the IPC error type (812). Following this, the error code of the IPC is set (814). Any IPC message for the source is added (816). Lastly, the blocked source of the IPC is set, so the IPC error is delivered to the proper task
10 (818).

20 Figure 12 describes how IPC error messages are delivered by the kernel (820). The kernel first determines if the monitor is on the IPC path source and destination (821). If so (822), the kernel determines if the source is blocked on the destination (823). If the source is blocked on the destination (824), then the kernel sets the IPC error state for the source (825) and unblocks the source (826). If the monitor is not on the proper IPC path (822) or the source is not blocked on the destination (824), the IPC error message is terminated and an error is returned to the monitor (410).

25 Figure 13 describes the mechanism for building unblock messages (910). The key step is the setting of the IPC unblock type (911). Following this, the error code of the IPC is set (912). Any IPC message for the source is added (816). Lastly, the blocked source of the IPC is set, so the IPC unblock is delivered to the proper task
30 (818).

5

10

15
20
25

Figure 14 describes how IPC unblock messages are delivered by the kernel (920). The kernel first sets the blocked source identity to that of the source in the IPC (921). Next, the kernel determines if the monitor is on the IPC path between the blocked source and the current destination (821). If so (822), the kernel determines if the source is blocked on the destination (823). If the source is blocked on the destination (824), then the kernel sets the IPC error state for the source to that specified in the unblock IPC (825) and unblocks the blocked source (1070). If the monitor is not on the proper IPC path (822) or the sender is not blocked on the destination (824), the IPC error message is terminated and an error is returned to the monitor (410).

Figure 15 describes how a IPC notify message is built that informs a controlling monitor of a completed IPC. First, the kernel obtains the blocked source of the active IPC (1050). If there is a blocked source (1060), the kernel sets the true source of the notify IPC to the blocked source (1032). If there is no blocked source (1060), then the true source of the notify IPC is set to the true source of the active IPC (1033). The kernel sets the destination in the notify IPC to the controlling monitor in the active IPC (1034). Lastly, the kernel sets the IPC type to notify (1035).

30
Figure 16 describes how the kernel delivers a notify IPC. The kernel first verifies that the controlling monitor of the active IPC is on the IPC path between the notify IPC's true source and the active IPC's destination (1041). If it is on the IPC path (822), then the kernel

delivers the notify IPC as a typical IPC (200). Otherwise, the monitor receives a notify error from the kernel in its error state (1042).

5 Figure 17 describes how the kernel updates the blocked source's destination when the monitor specifies a change in destination (770). First, the kernel gets the blocked source value (1050). If there is a blocked source (1060), then the source of the changed destination is then set to the blocked source (771). Otherwise, the source of the changed destination is set to the true source (772). Next, the new destination is retrieved (823). If the source is blocked (824), then the destination in which it is blocked is changed to the new destination (775). Otherwise, no action is taken (774).

10

15

20 While it is apparent that the invention herein disclosed is well calculated to fulfill the objects stated above, it will be appreciated that numerous modification and embodiments may be devised by those skilled in the art, and it is intended that the appended claims cover all such modifications and embodiments as fall within the true spirit and scope of the present invention.

25

1 CLAIMS

2 1. A method for controlling the synchrony of interprocess
3 communication (IPC) which uses: a. the original sender of
4 the communication, called the source; b. the ultimate
5 receiver of the communication, called the destination; c.
6 a kernel process that is capable of delivering
7 communications and unblocking the source of the
8 communication; and d. one or more processes, called
9 monitors, to which the kernel process may choose to
10 deliver a communication to, rather than the destination,
11 and which have the ability to request control of when a
12 source is unblocked, called the controlling monitor of
13 the source. wherein the method for controlling the
14 synchrony of interprocess communication comprise the
15 following steps:
16 d. when the communication is sent by the source, the
17 source is blocked by the kernel process;
18 e. the kernel either delivers the communication to the
19 destination or to a monitor;
20 f. the monitor may request that the kernel process make
21 it the controlling monitor of the source;
22 g. the monitor forwards the communication to the
23 destination specifying the source of the communication,
24 and the monitor is blocked by the kernel process until
25 the communication is delivered either to the destination
26 or another monitor;
27 h. when the communication is delivered to the
28 destination, the kernel process either: (1) signals the
29 controlling monitor of the specified source, if any, or
 (2) unblocks the specified source;

30 i. the controlling monitor of the source can request that
31 the kernel process unblock the source at any time after
32 being signalled that the communication was delivered to
33 the destination.

1 2. A method as claimed in claim 1, wherein the monitor
2 may change the identity of the source that is seen by the
3 destination while maintaining the identity of the
4 original source for unblocking purposes.

1 3. A method as claimed in claim 1, wherein the monitor
2 may change the identity of the destination to which the
3 communication is to be delivered.

① 4. A method as claimed in claim 1, wherein the message
② embodied in the communication is not delivered to the
③ monitor.

④ 5. A method as claimed in claim 1, wherein the timeouts
⑤ on the communication are interpreted as being relative to
⑥ the behaviors of the source and destination.

⑦ 6. A method as claimed in claim 5, wherein the kernel
⑧ process does not deliver the communication to the
⑨ destination or a monitor until the destination is ready
⑩ to receive the communication and the destination timeout
⑪ (i.e., the timeout set by the source on the amount of
⑫ time it will wait for the destination to become ready to
⑬ receive the communication) has not expired.

1 7. A method as claimed in claim 5, wherein the kernel
2 process delivers the communication to the monitor as soon

3 as the monitor is ready, and the monitor checks for the
4 destination timeout's expiration.

1 8. A method as claimed in claim 5, wherein the kernel
2 process verifies the source timeout (i.e., the timeout
3 set by the destination on the amount of time it will wait
4 for the source to initiate the communication) has not
5 expired before sending the communication to the
6 destination or a monitor.

1 9. A method as claimed in claim 1, wherein multiple
2 monitor processes may claim to be the controlling monitor
3 of a source.

10. A method as claimed in claim 1, wherein the kernel
process may authorize a monitor's permission to be the
controlling monitor of a particular source.

11. A method as claimed in claim 10, wherein the kernel
process may authorize a monitor's permission to be the
controlling monitor of any source.

12. A method as claimed in claim 1, wherein the
controlling monitor for a particular source may be stored
in the kernel.

13. A method as claimed in claim 9, wherein a sequence of
controlling monitors for a particular source may be
stored in the kernel.

1 14. A method as claimed in claim 1, wherein the identity
2 of the controlling monitor of a particular source is
3 passed in the IPC to the destination.

1 15. A method as claimed in claim 14, wherein the sequence
2 of controlling monitors for a particular source is stored
3 by the controlling monitor storing its controlling
4 monitor predecessor for each source.

1 16. A method as claimed in claim 1, wherein the
2 controlling monitor for a source is implemented by
3 changing the original source to the controlling monitor
4 and the controlling monitor stores the identity of the
5 original source.

1 17. A method as claimed in claim 16, wherein a sequence
2 of controlling monitors for a particular source is
3 implemented as a sequence of original source changes in
4 the monitors where the last is the true original source.

1 18. A method as claimed in claim 1, wherein the monitors
2 are implemented as threads in the same process.

1 19. A method as claimed in claim 1, wherein the monitors
2 are implemented as procedures in the same process.

1 20. A method as claimed in claim 1, wherein the monitor
2 procedures are in the kernel process.

1 21. A method of performing interprocess communications
2 (IPCs), comprising the steps of:

3 an IPC process receiving IPC requests, each of
4 the IPC requests including a source identifier
5 identifying a source and a destination identifier
6 identifying a destination;
7 building IPCs in response to the requests;
8 transmitting the IPCs from the sources to the
9 destinations;
10 intercepting and examining selected ones of the
11 IPCs; and
12 controlling the synchrony of the IPCs so that
13 each IPC appear to the source and to the destination to
14 be implemented according to the same semantics regardless
15 of whether the IPC is intercepted and examined.

22. A method according to Claim 21, wherein:

the step of building and transmitting IPCs includes the step of using a kernel to build and transmit the IPCs;

the step of intercepting and examining selected ones of the IPCs includes the step of using monitors to intercept and examine the selected ones of the IPCs; and

the controlling step includes the step of using the monitors as extensions of the kernel so that the IPCs appear to the sources and to the destinations to be implemented according to the same semantics regardless of whether a monitor is used or not used to intercept and examine the IPCs.

23. A method according to claim 22, wherein, for each IPC, one or more of the monitors manages the identity of the source for the IPC.

1 24. A method according to Claim 22, further comprising
2 the step of, for each IPC, blocking the source for the
3 IPC at selected times, and wherein one or more of the
4 monitors manages when the source is unblocked.

1 25. A method according to Claim 22, wherein, for each
2 IPC, one of the monitors manages whether the monitor is
3 notified when the IPC reaches the destination for the
4 IPC.

1 26. An interprocess communications (IPCs) system,
2 comprising:

3 a processor having an IPC process for receiving
4 IPC requests, each of the IPC requests including a source
5 identifier identifying a source for the IPC and a
6 destination identifier identifying a destination for the
7 IPC;

8 means for building IPCs in response to the
9 requests;

10 means for transmitting the IPCs from the
11 sources to the destinations;

12 means for intercepting and examining selected
13 ones of the IPCs; and

14 a controller for controlling the synchrony of
15 the IPCs so that each IPC appear to the source and to the
16 destination for the IPC to be implemented according to
17 the same semantics regardless of whether the IPC is
18 intercepted and examined.

1 27. A system according to Claim 26, wherein:

2 the means for building and transmitting IPCs
3 includes a kernel to build and transmit the IPCs;

4 the means for intercepting and examining
5 selected ones of the IPCs include monitors to intercept
6 and examine the selected ones of the IPCs; and

7 the monitors are used as extensions of the
8 kernel so that the IPCs appear to the sources and to the
9 destinations to be implemented according to the same
10 semantics regardless of whether a monitor is used or not
11 used to intercept and examine the IPCs.

1 28. A system according to claim 26, wherein, for each
2 IPC, one or more of the monitors manages the identity of
3 the source for the IPC.

1 29. A system according to Claim 26, wherein, for each
2 IPC, the source for the IPC is blocked at selected times,
3 and one or more of the monitors manages when the source
4 is unblocked.

1 30. A system according to Claim 26, wherein, for each
2 IPC, one of the monitors manages whether the monitor is
3 notified when the IPC reaches the destination for the
4 IPC.

1 31. A program storage device readable by machine,
2 tangibly embodying a program of instructions executable
3 by the machine to perform method steps for performing
4 interprocess communications (IPCs), said method steps
5 comprising:

6 an IPC process receiving IPC requests, each of
7 the IPC requests including a source identifier
8 identifying a source and a destination identifier
9 identifying a destination;

10 building IPCs in response to the requests;
11 transmitting the IPCs from the sources to the
12 destinations;
13 intercepting and examining selected ones of the
14 IPCs; and
15 controlling the synchrony of the IPCs so that
16 each IPC appear to the source and to the destination to
17 be implemented according to the same semantics regardless
18 of whether the IPC is intercepted and examined.

1 32. A program storage device according to Claim 31,
2 wherein:

3 the step of building and transmitting IPCs
4 includes the step of using a kernel to build and transmit
5 the IPCs;

6 the step of intercepting and examining selected
7 ones of the IPCs includes the step of using monitors to
8 intercept and examine the selected ones of the IPCs; and

9 the controlling step includes the step of using
10 the monitors as extensions of the kernel so that the IPCs
11 appear to the sources and to the destinations to be
12 implemented according to the same semantics regardless of
13 whether a monitor is used or not used to intercept and
14 examine the IPCs.

1 33. A program storage device according to claim 32,
2 wherein, for each IPC, one or more of the monitors
3 manages the identity of the source for the IPC.

1 34. A program storage device according to Claim 32,
2 further comprising the step of, for each IPC, blocking
3 the source for the IPC at selected times, and wherein one

4 or more of the monitors manages when the source is
5 unblocked.

1 35. A program storage device according to Claim 32,
2 wherein, for each IPC, one of the monitors manages
3 whether the monitor is notified when the IPC reaches the
4 destination for the IPC.

00000000000000000000000000000000

A MECHANISM FOR SYNCHRONOUS INTERPROCESS
COMMUNICATION OVER TRANSPARENT EXTERNAL MONITORS

ABSTRACT

10 A method and system for performing interprocess communications (IPCs). The method comprises the steps of receiving IPC requests, where each of the IPC requests identifies a source and a destination; building IPCs in response to the request; transmitting the IPCs from the sources to the destinations; and intercepting and examining selected ones of the IPCs. The method comprises the further step of controlling the synchrony of the IPCs so that each IPC appears to its source and destination to be implemented according to the same semantics regardless of whether the IPC is intercepted and examined. With the preferred embodiment of this invention, the system monitors are considered as an extension of the system kernel (although they may be linked into the kernel and run in kernel mode as well), so the source and destination are treated as if the kernel is still processing the IPC. Thus, the desired semantics of communication can be implemented in the monitors.

15

20

25

Synchronous IPC over Monitors

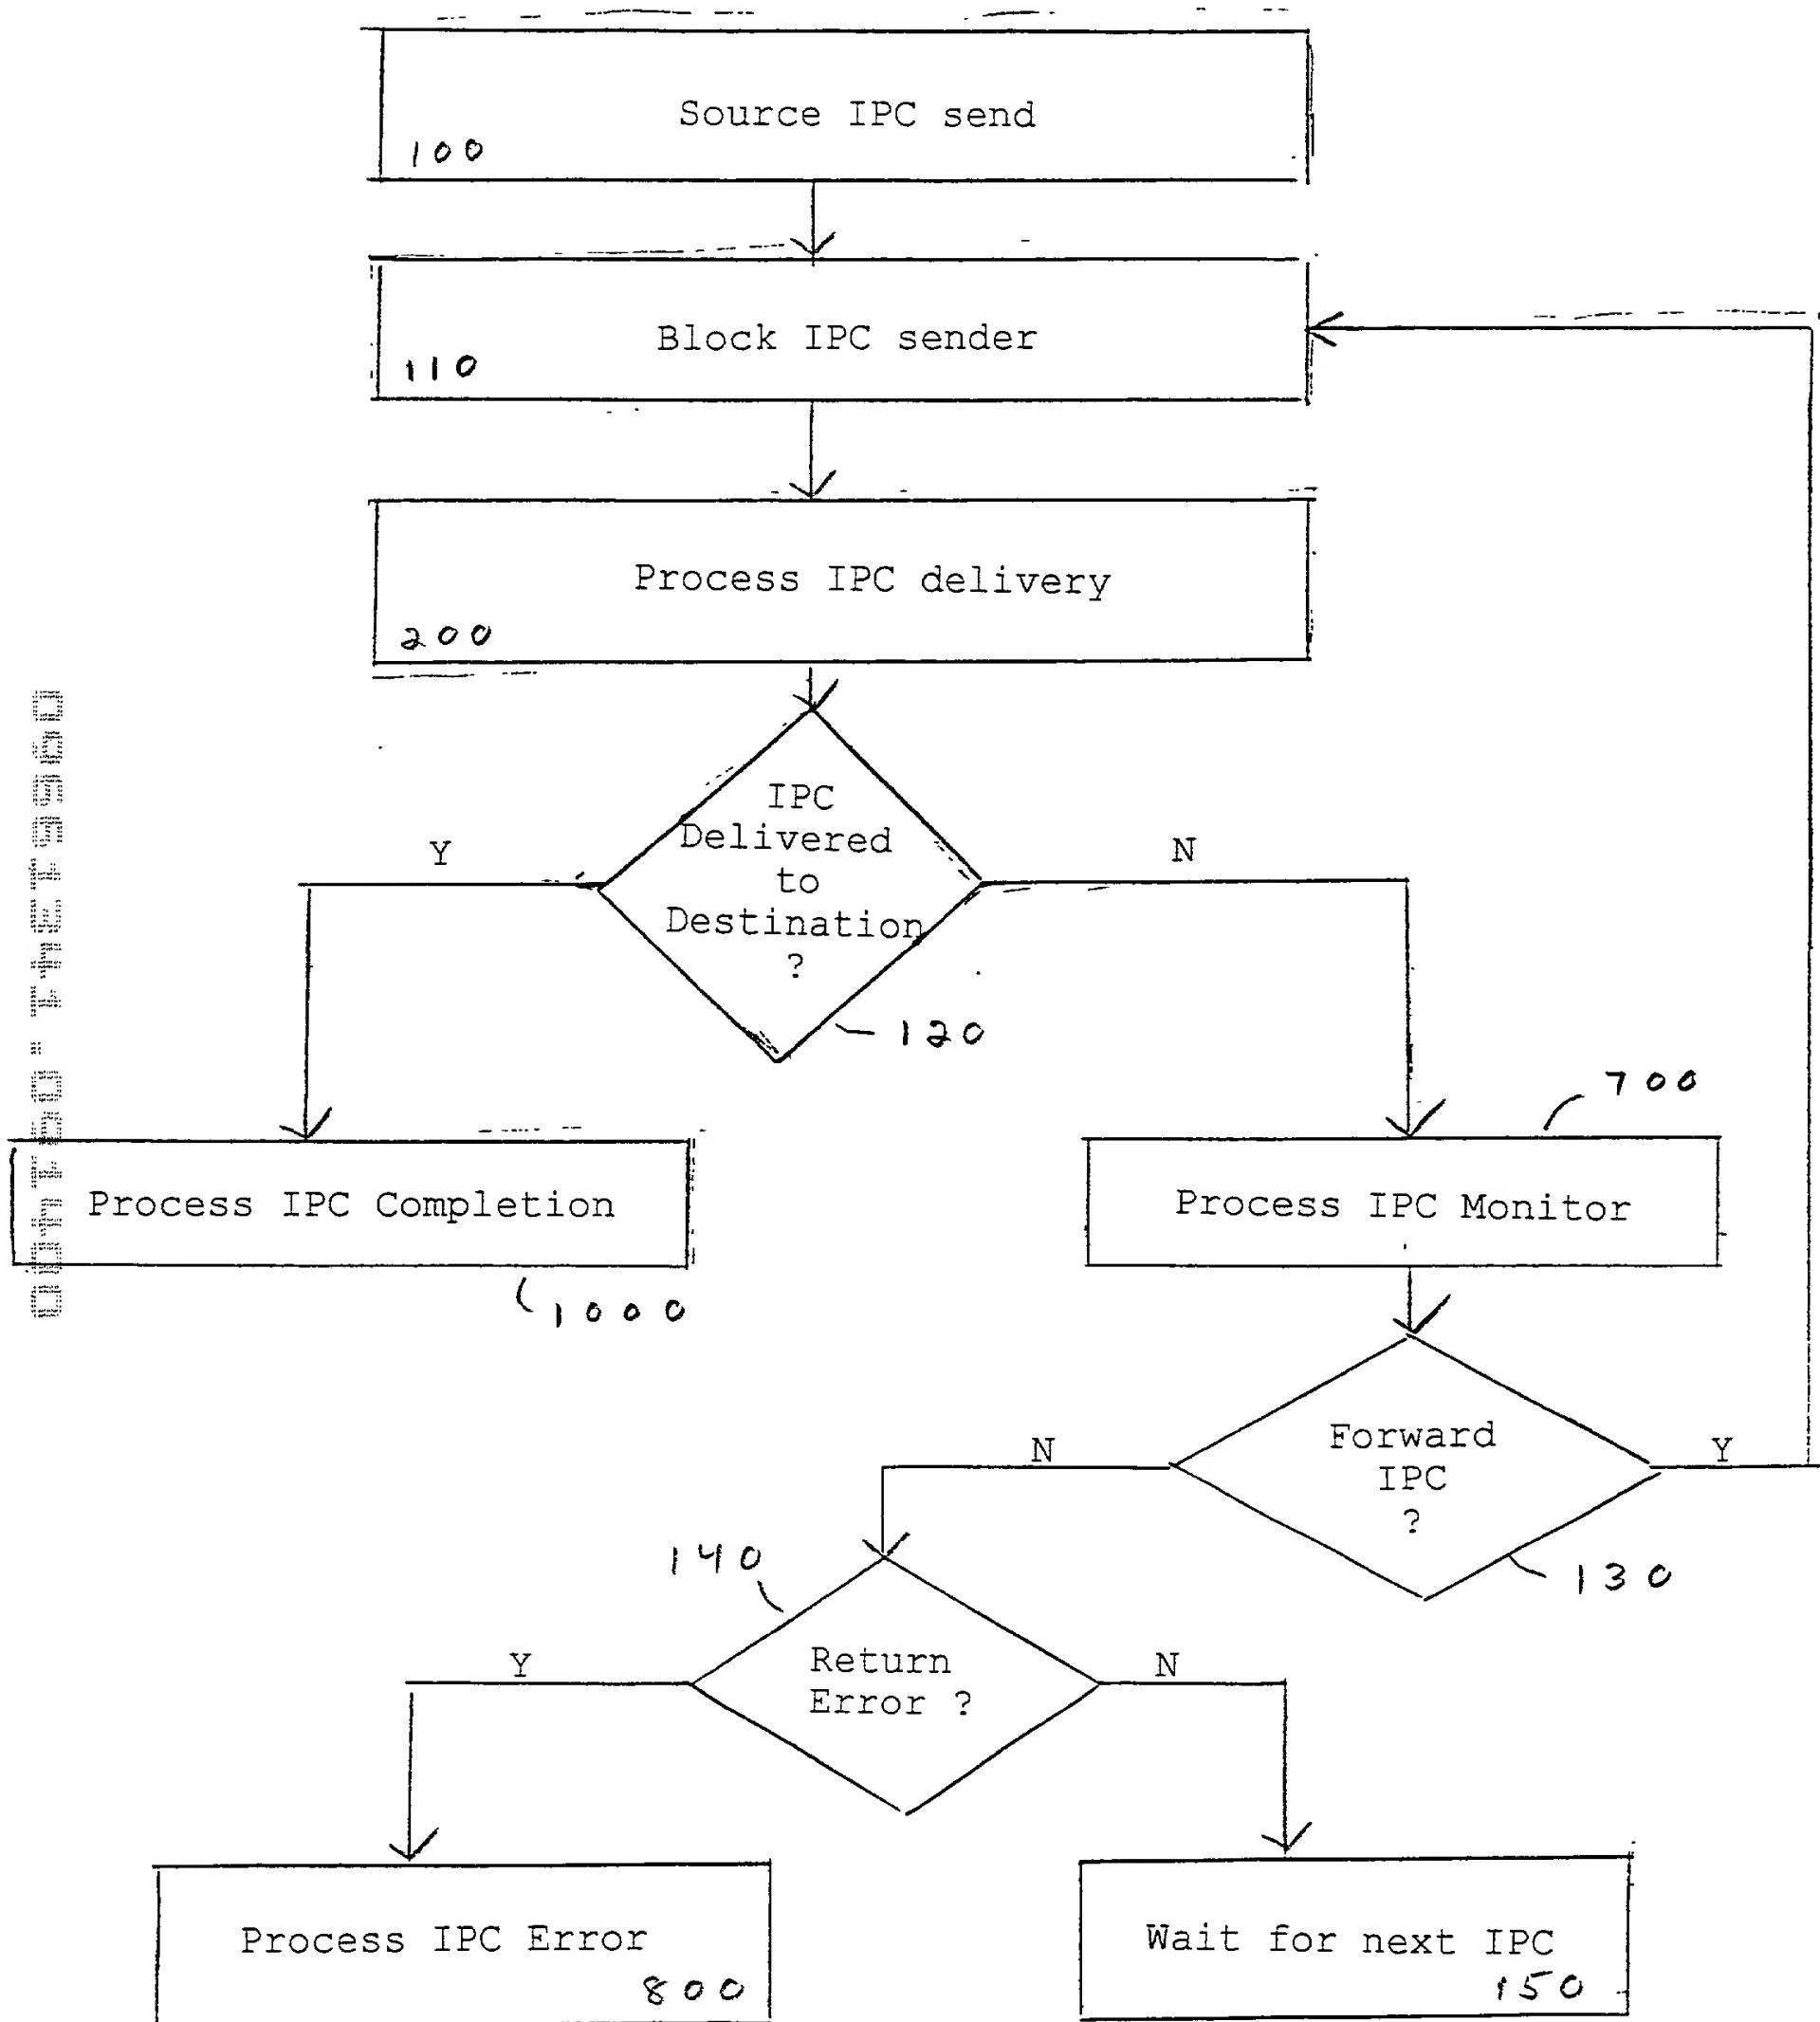


Fig. 1

Process IPC Delivery

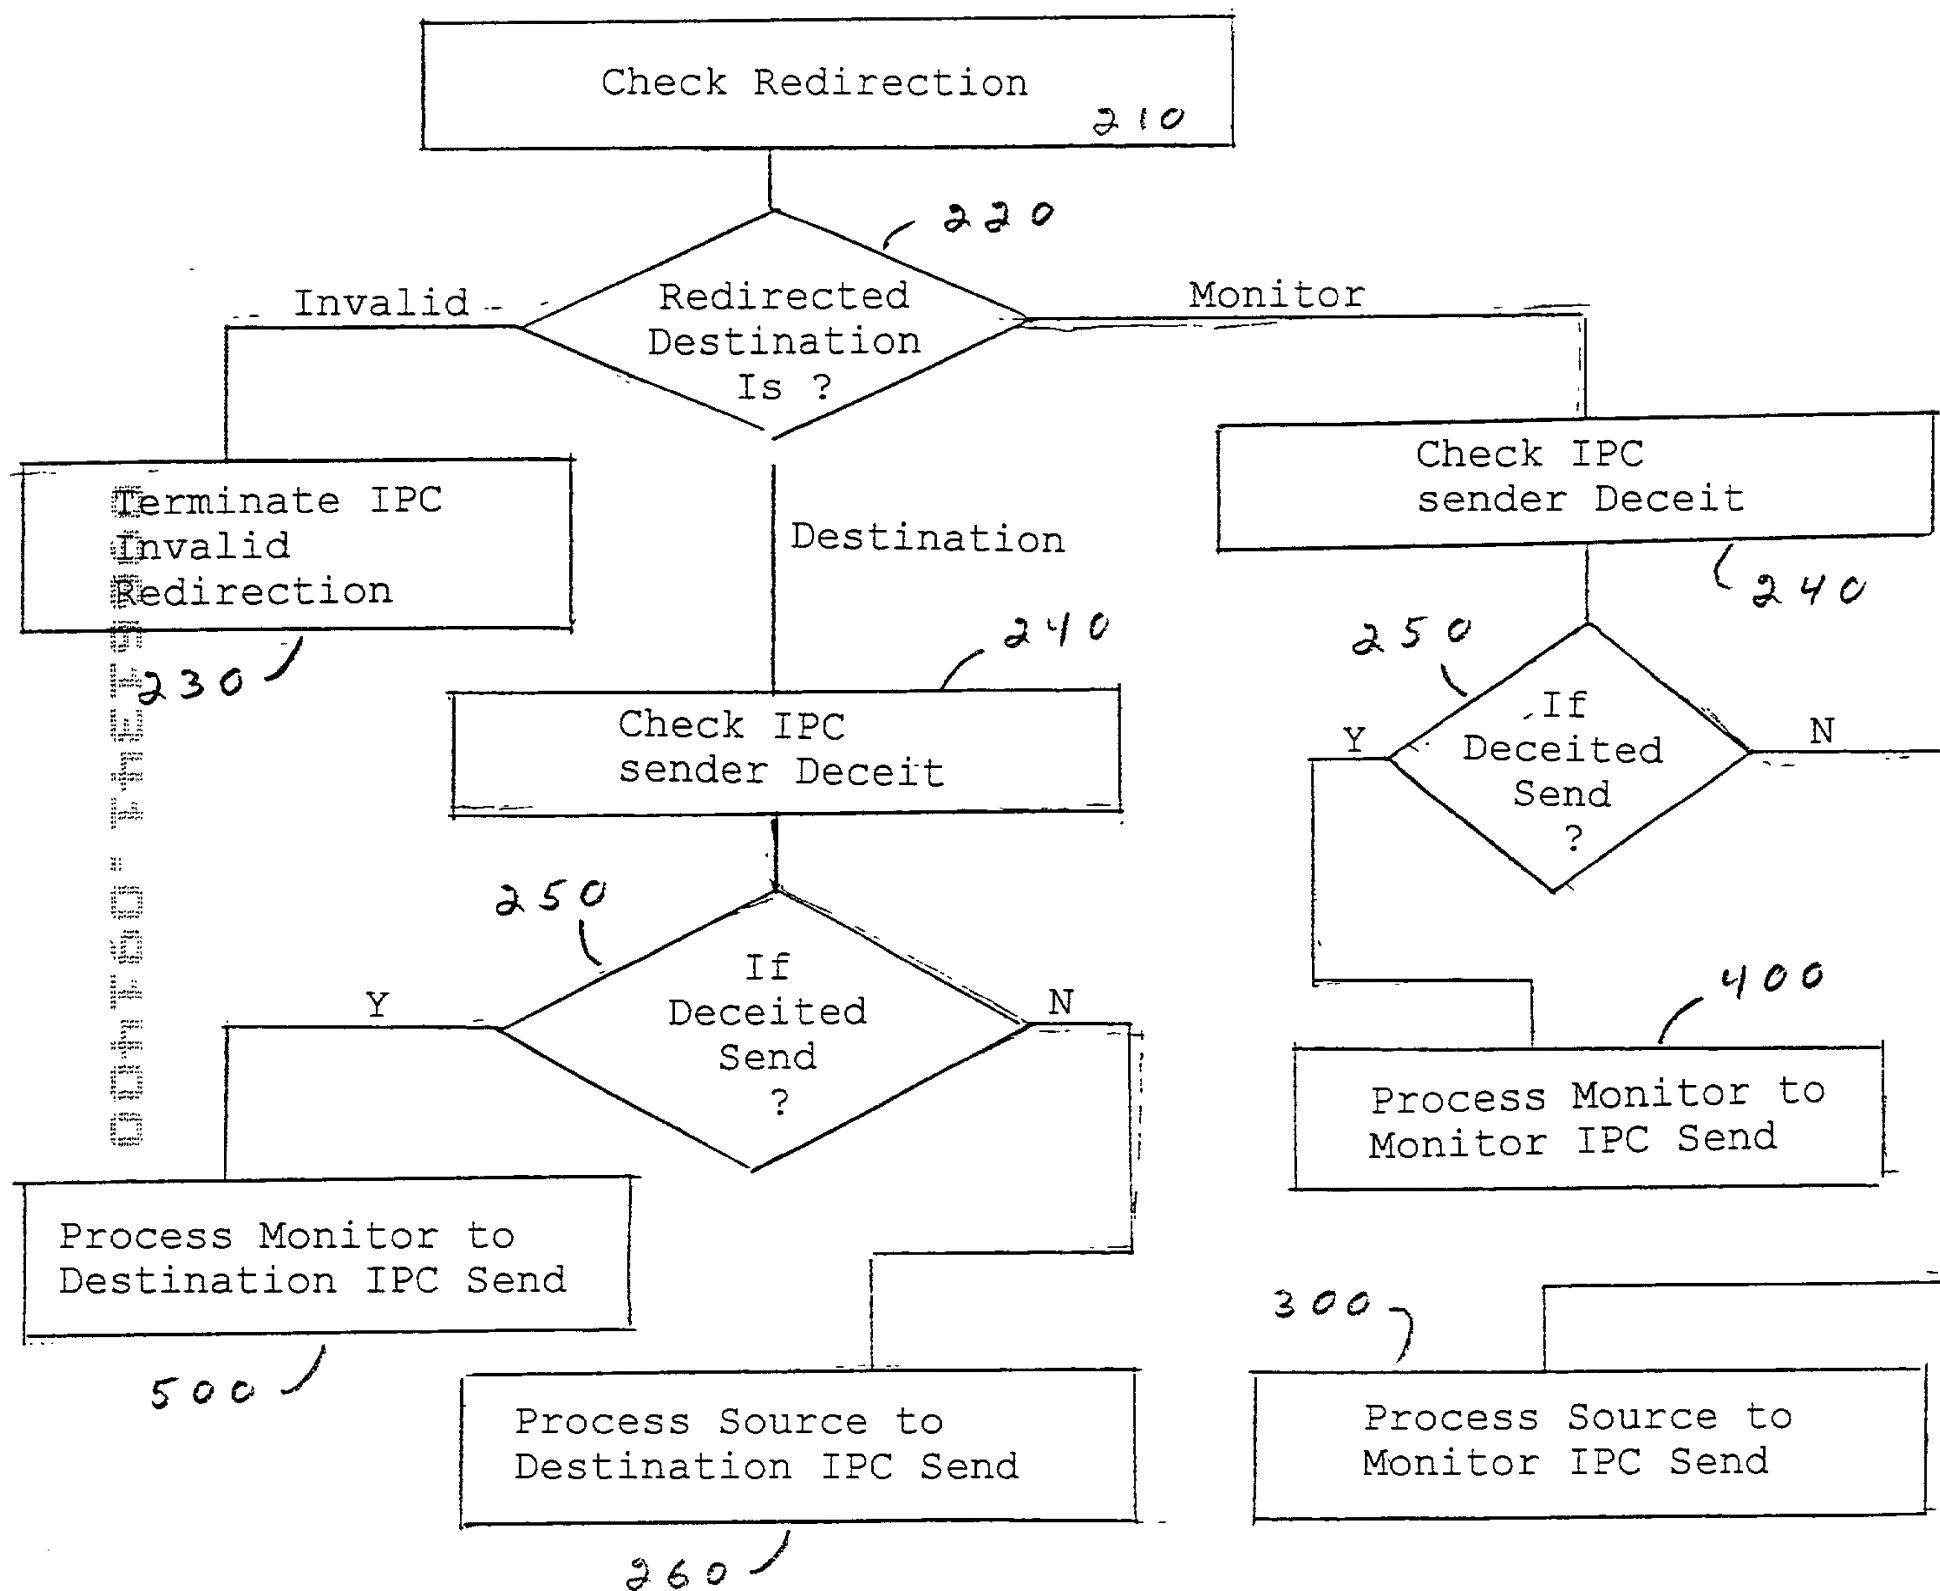


Fig. 2

3/15
YOR920000575US1

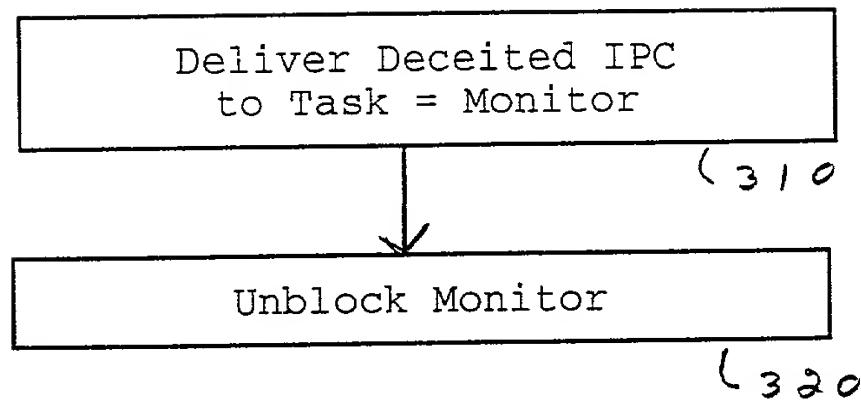


Fig. 3

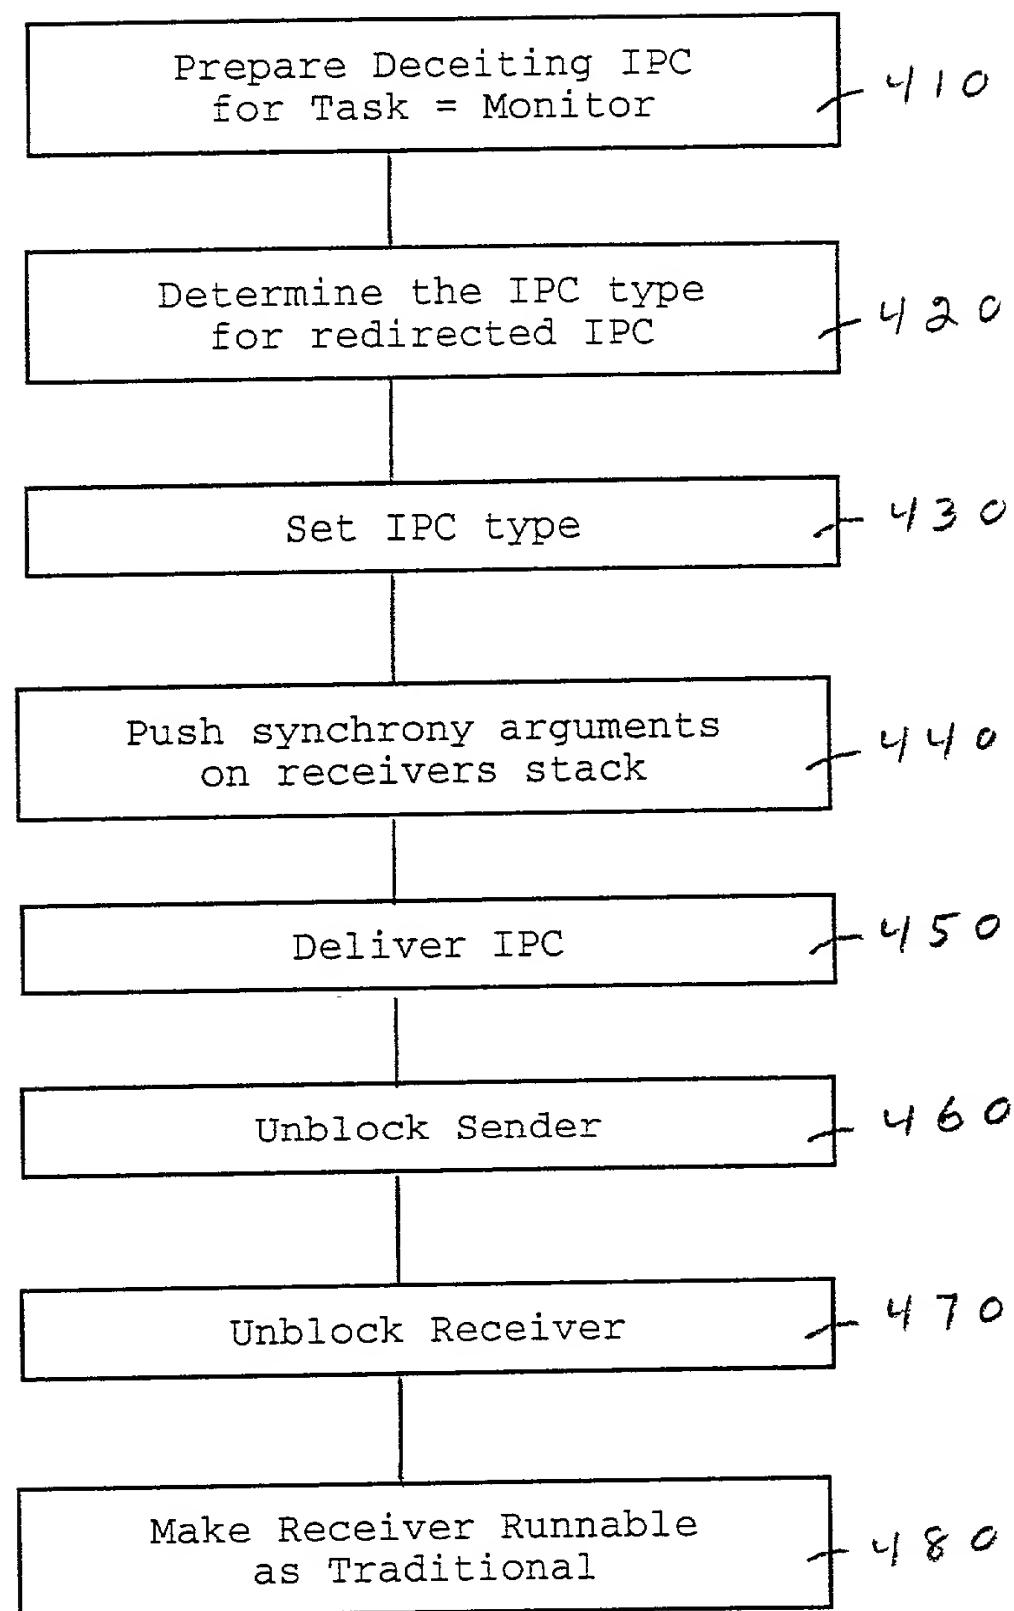


Fig. 4

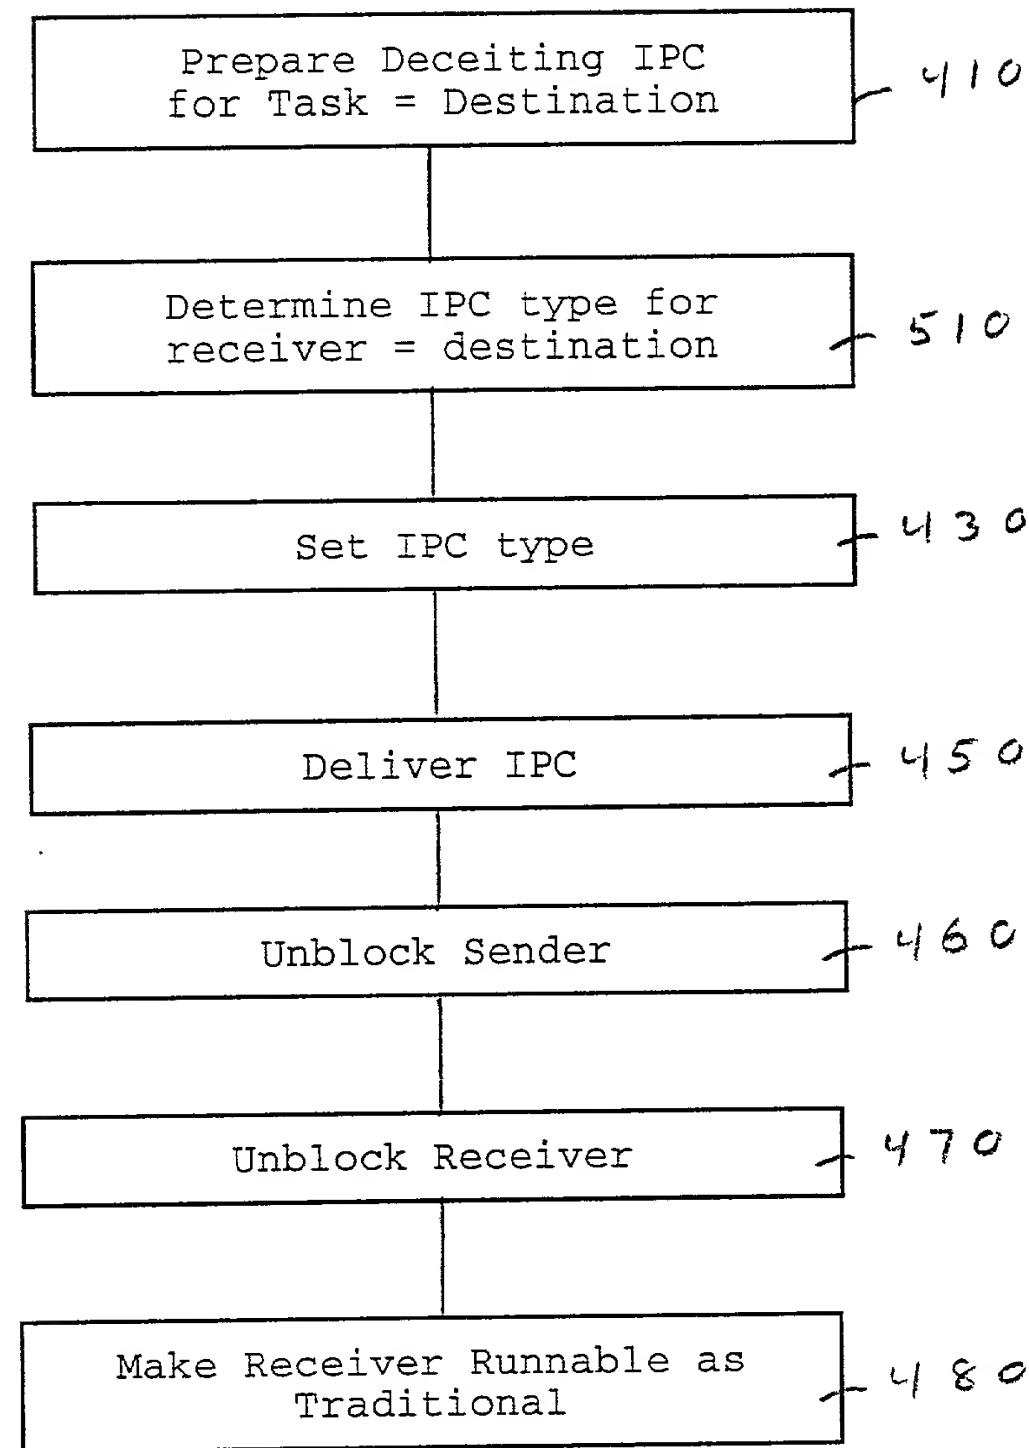
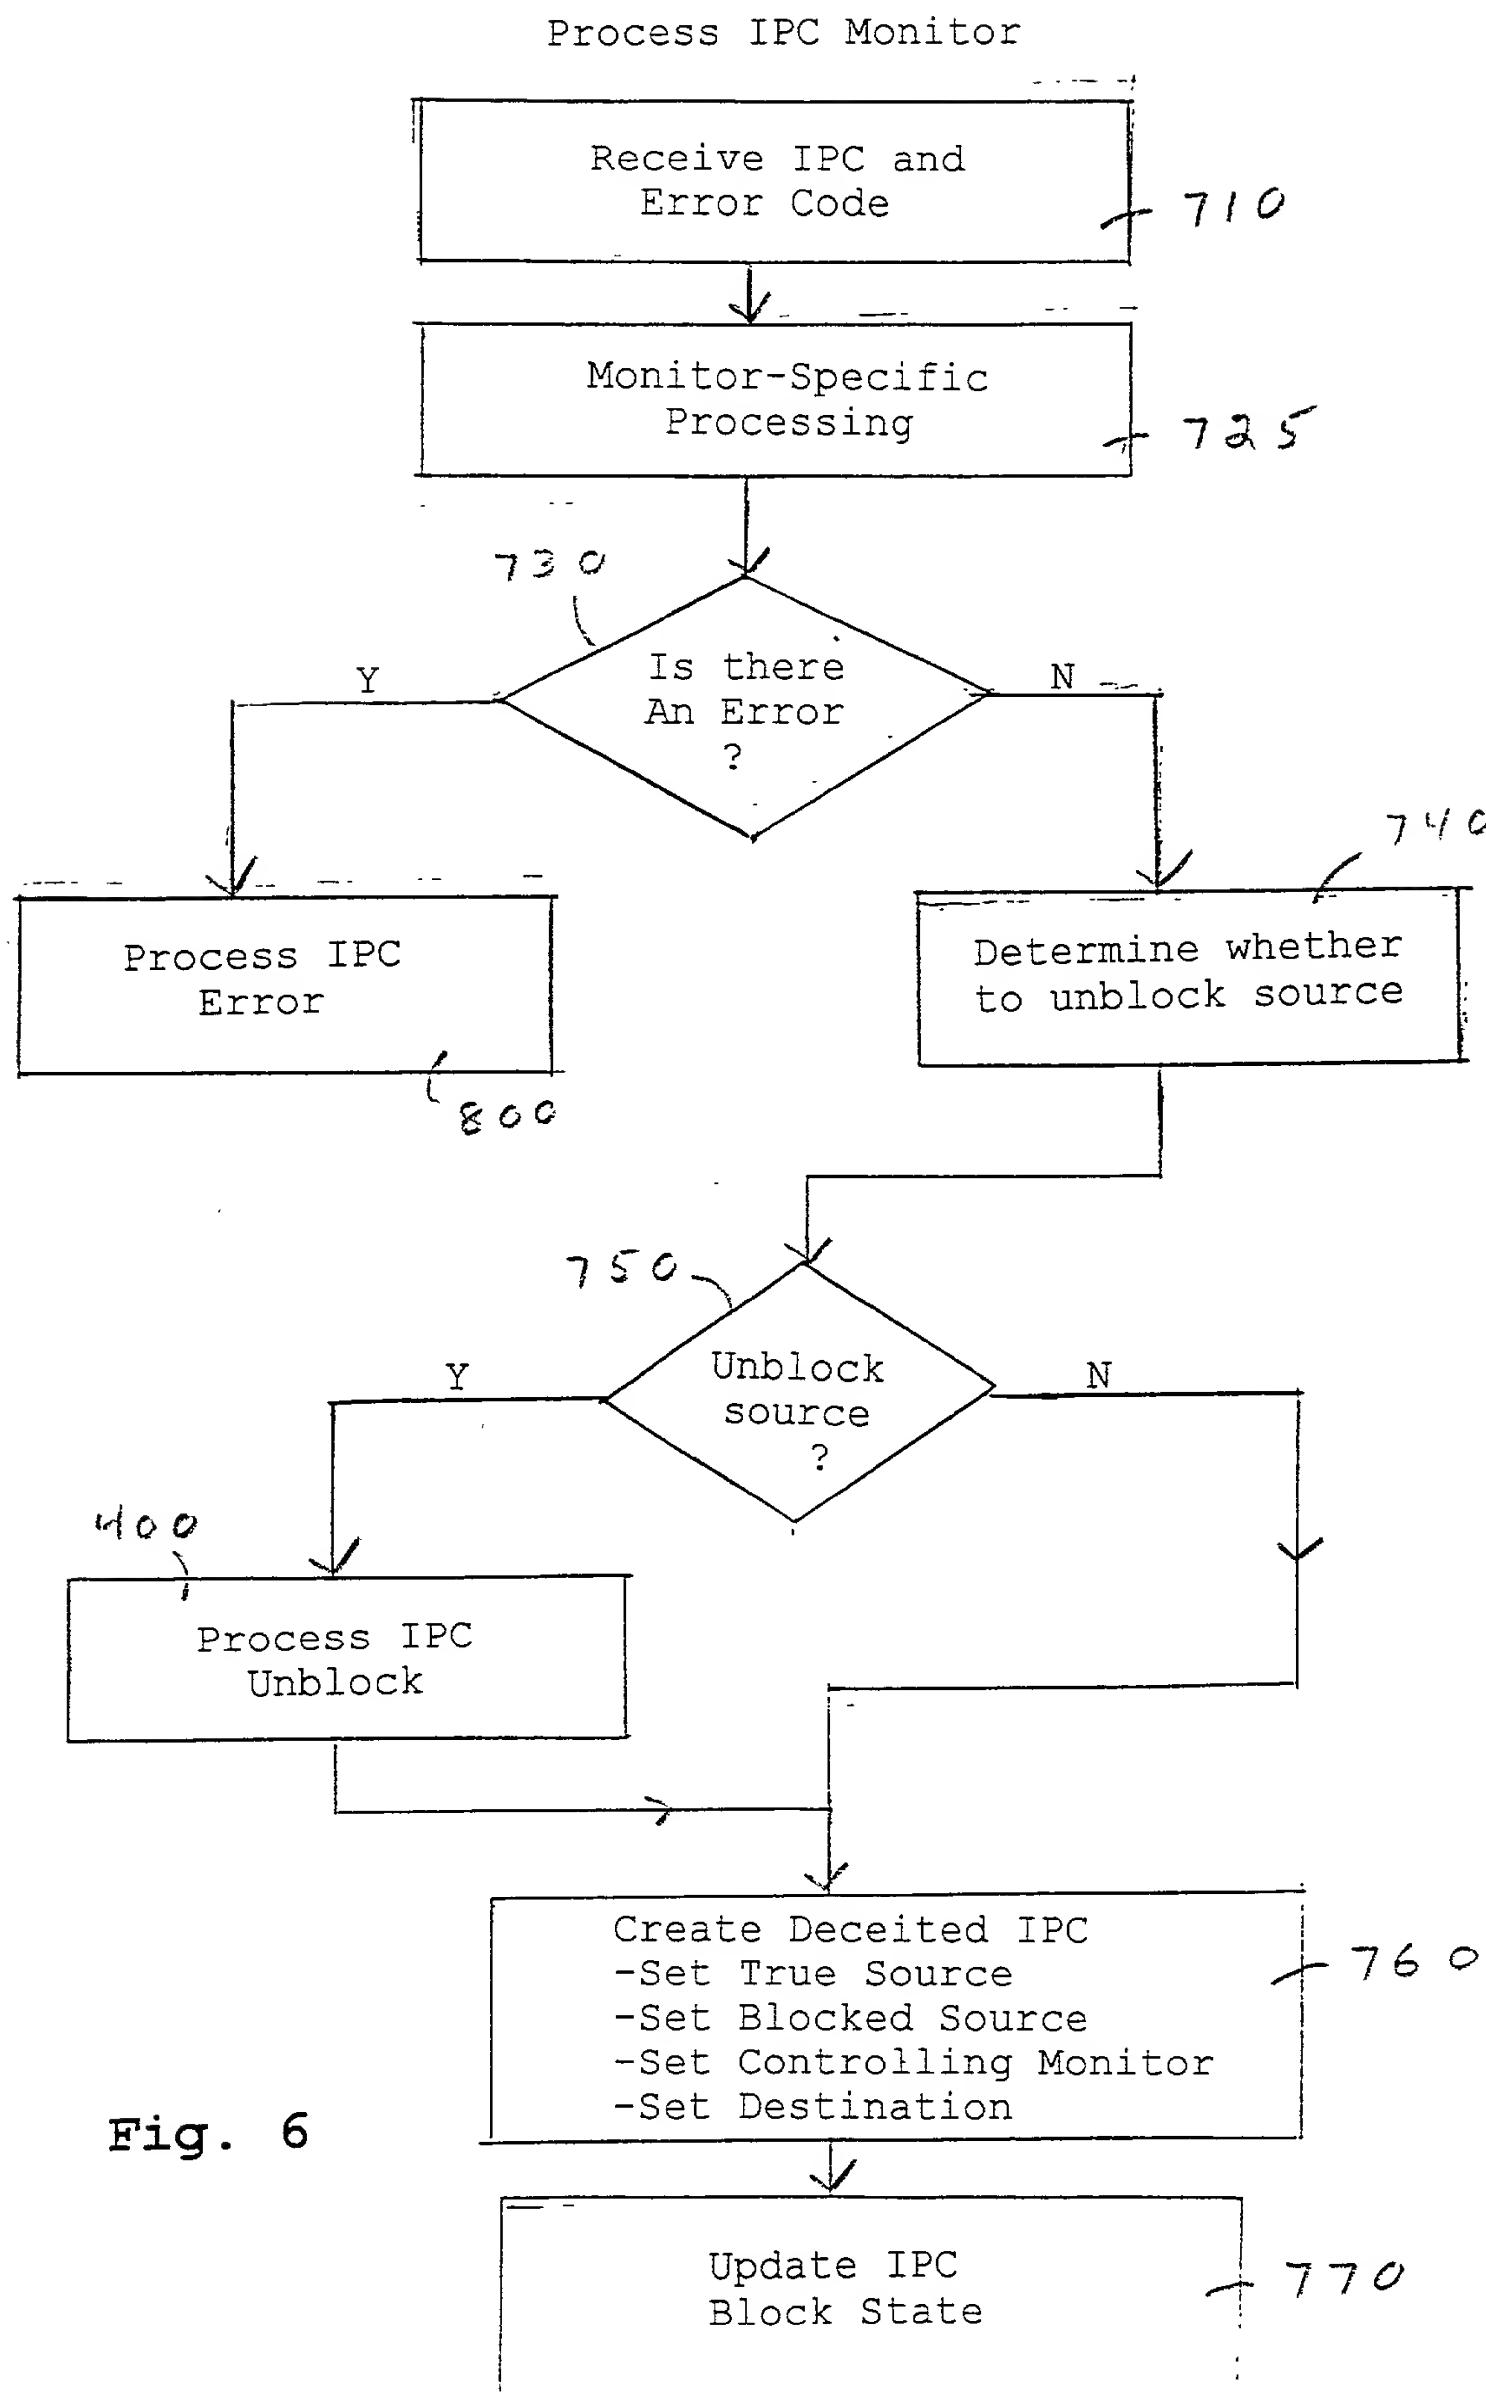


Fig. 5



7/15
YOR920000575US1

Fig. 7

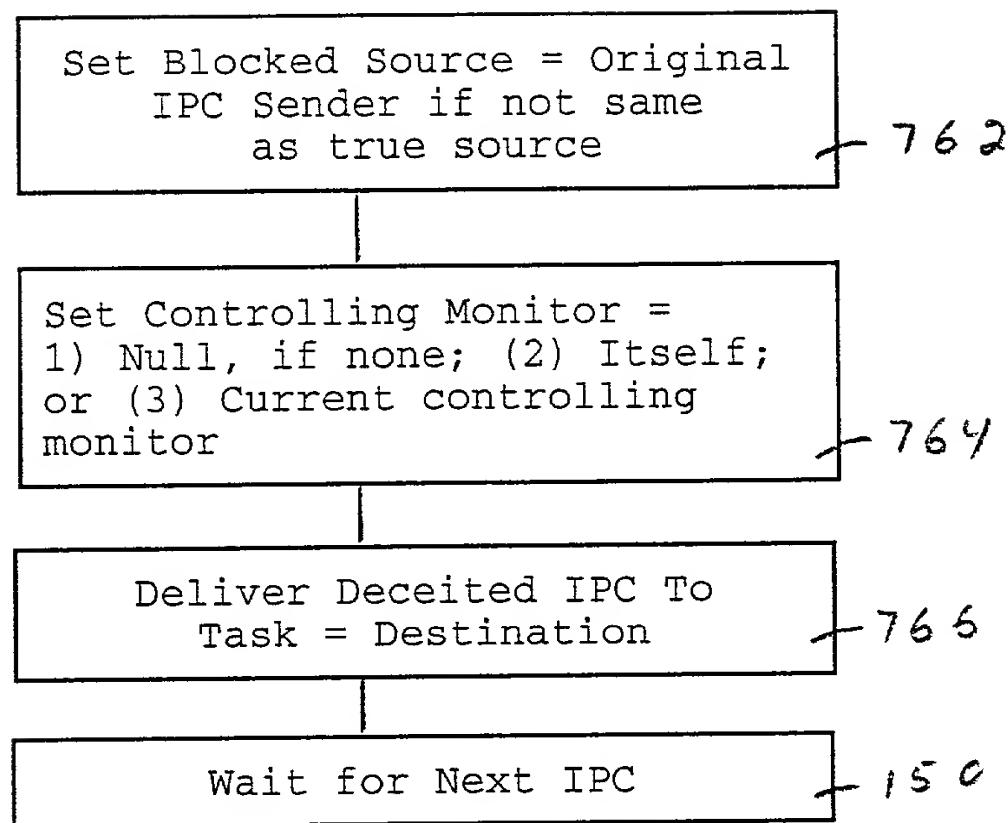


Fig. 8

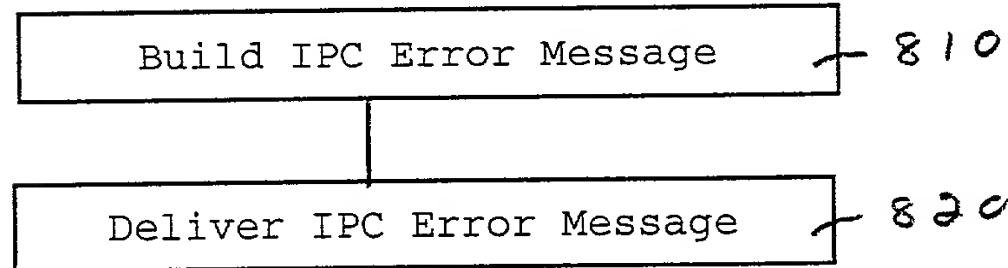
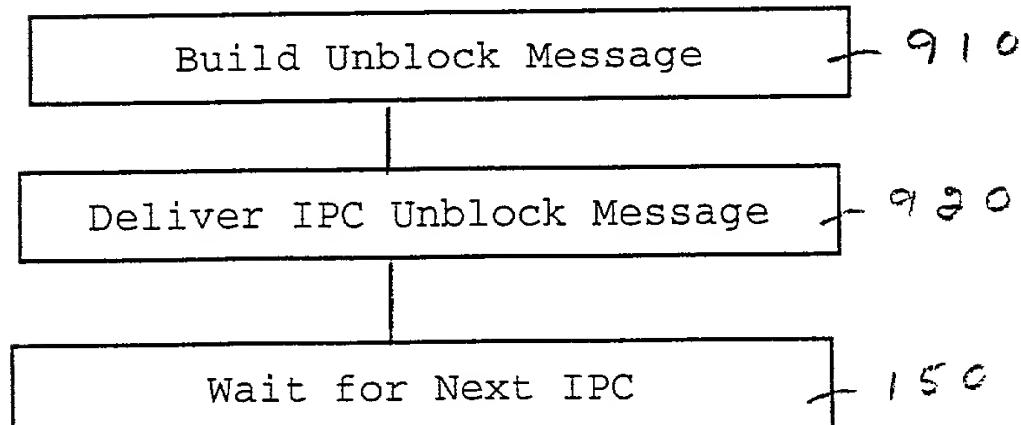


Fig. 9



Process IPC Computation

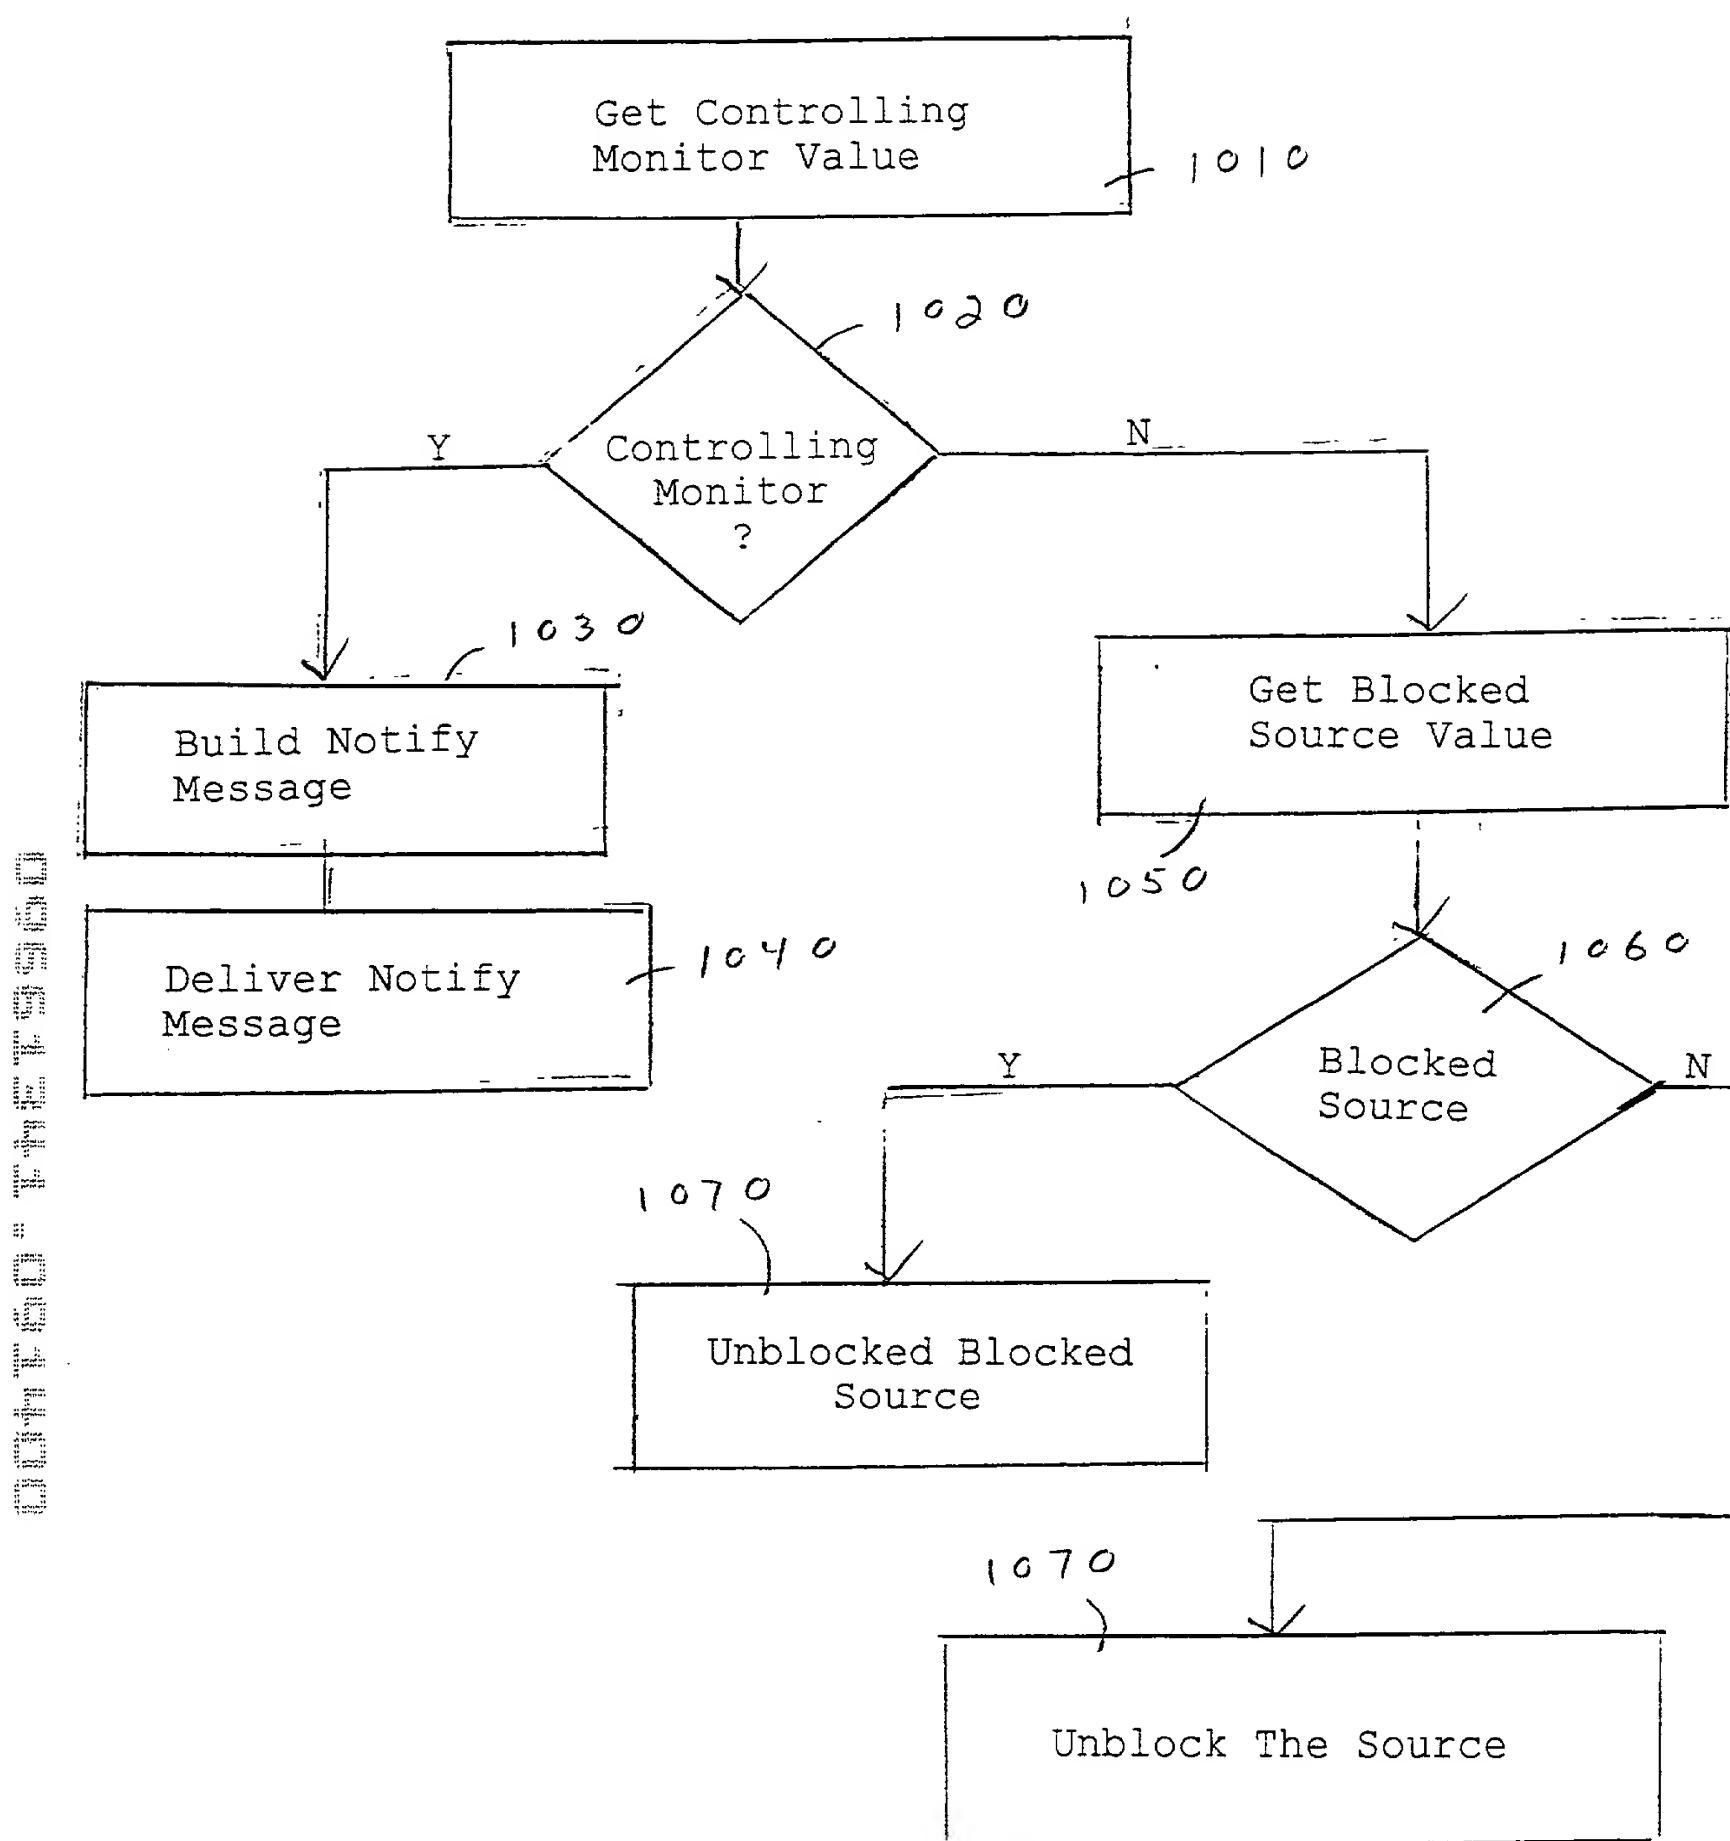


Fig. 10

9/15
YOR920000575US1

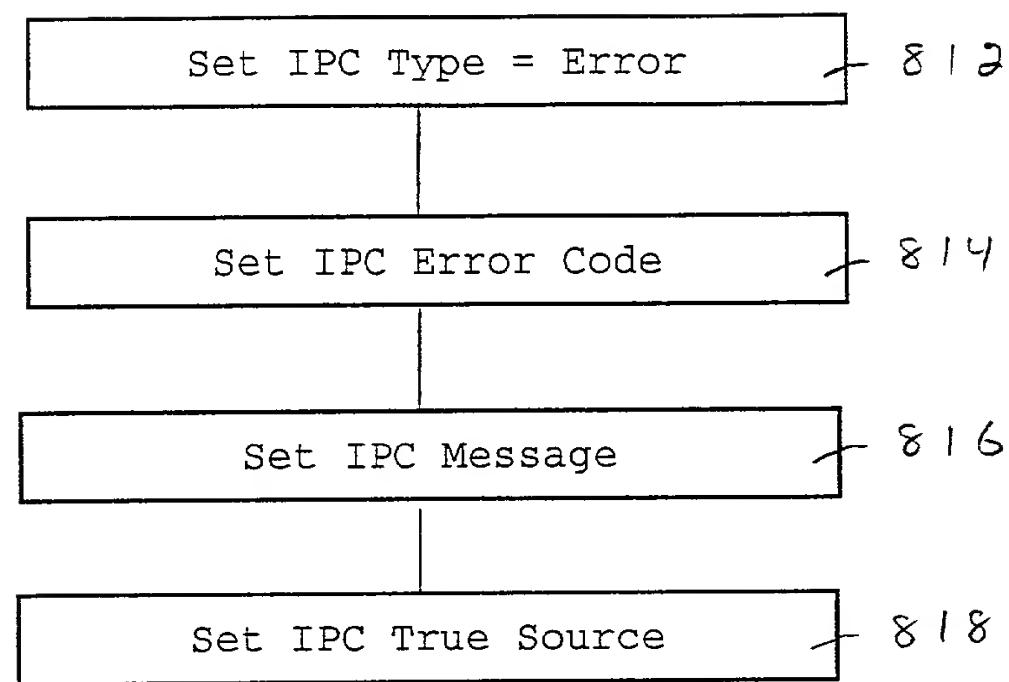
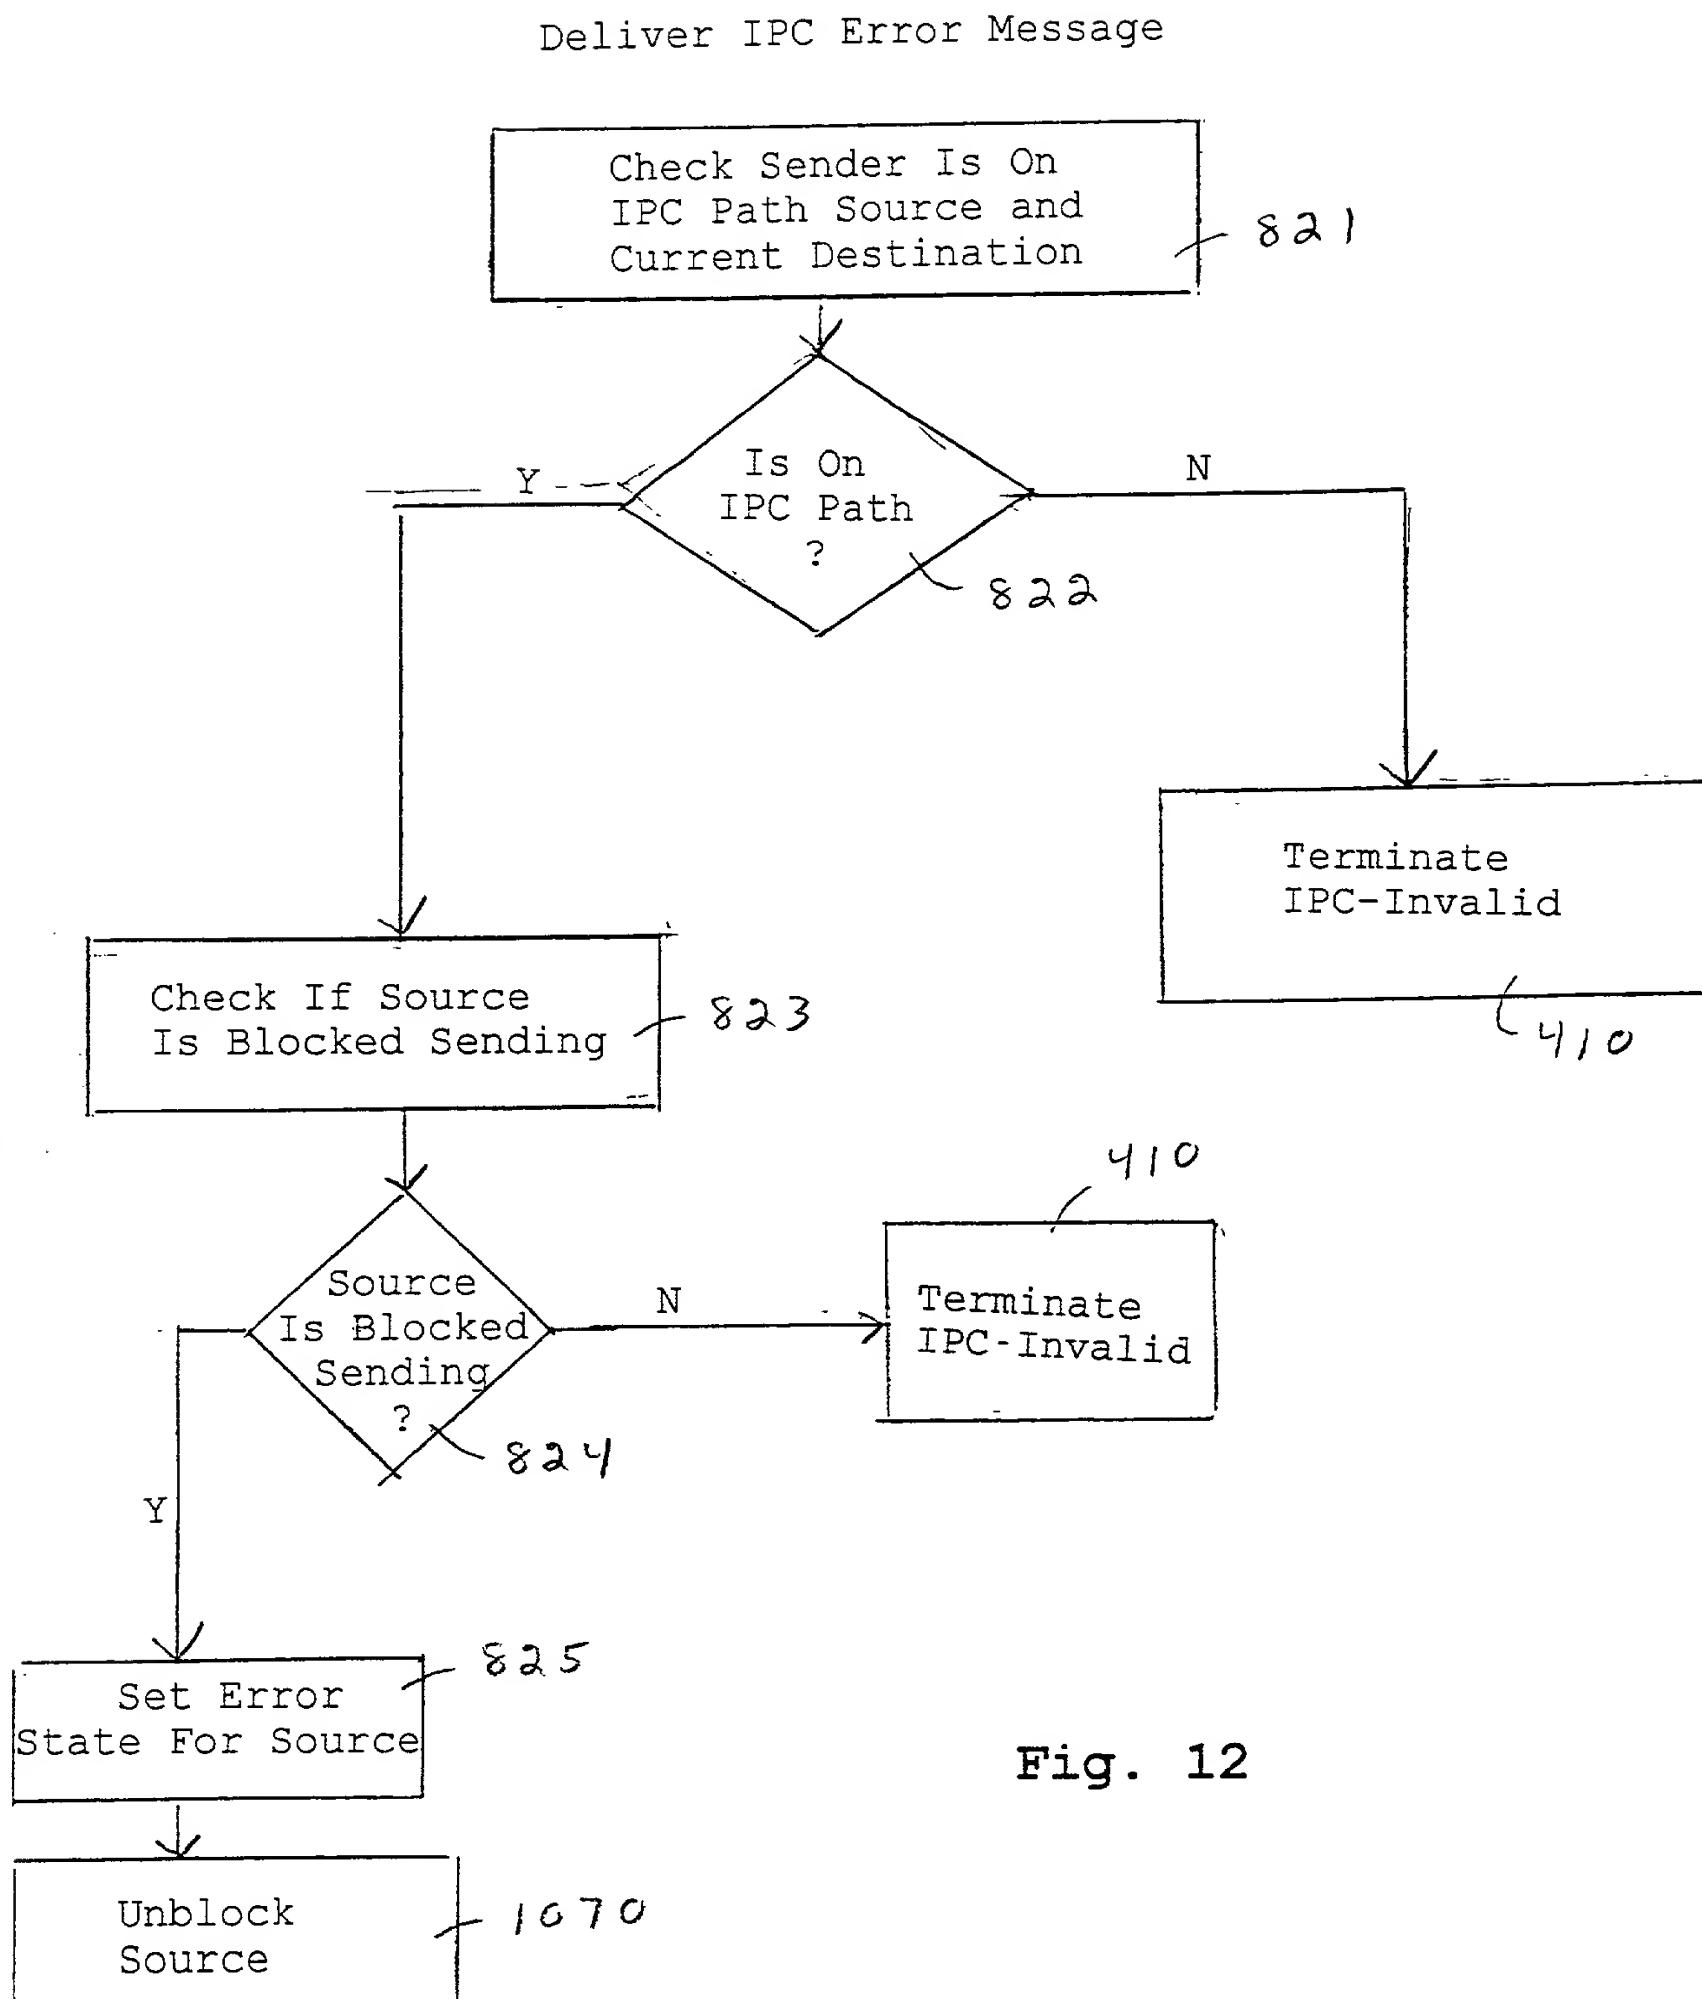


Fig. 11



11/15
YOR920000575US1

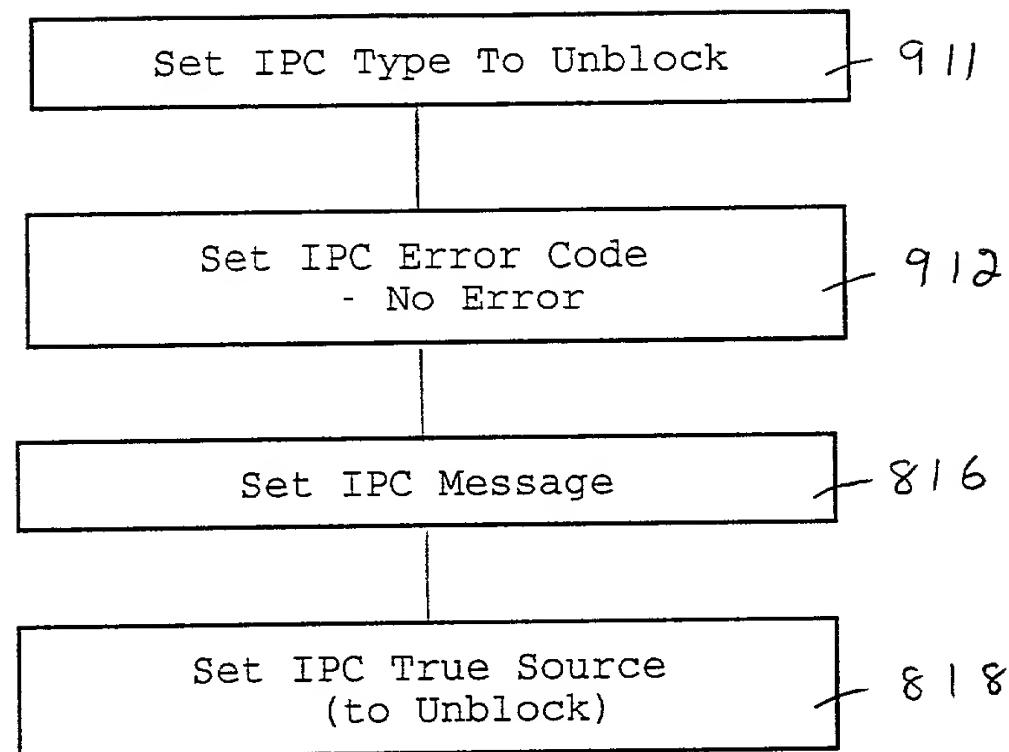


Fig. 13

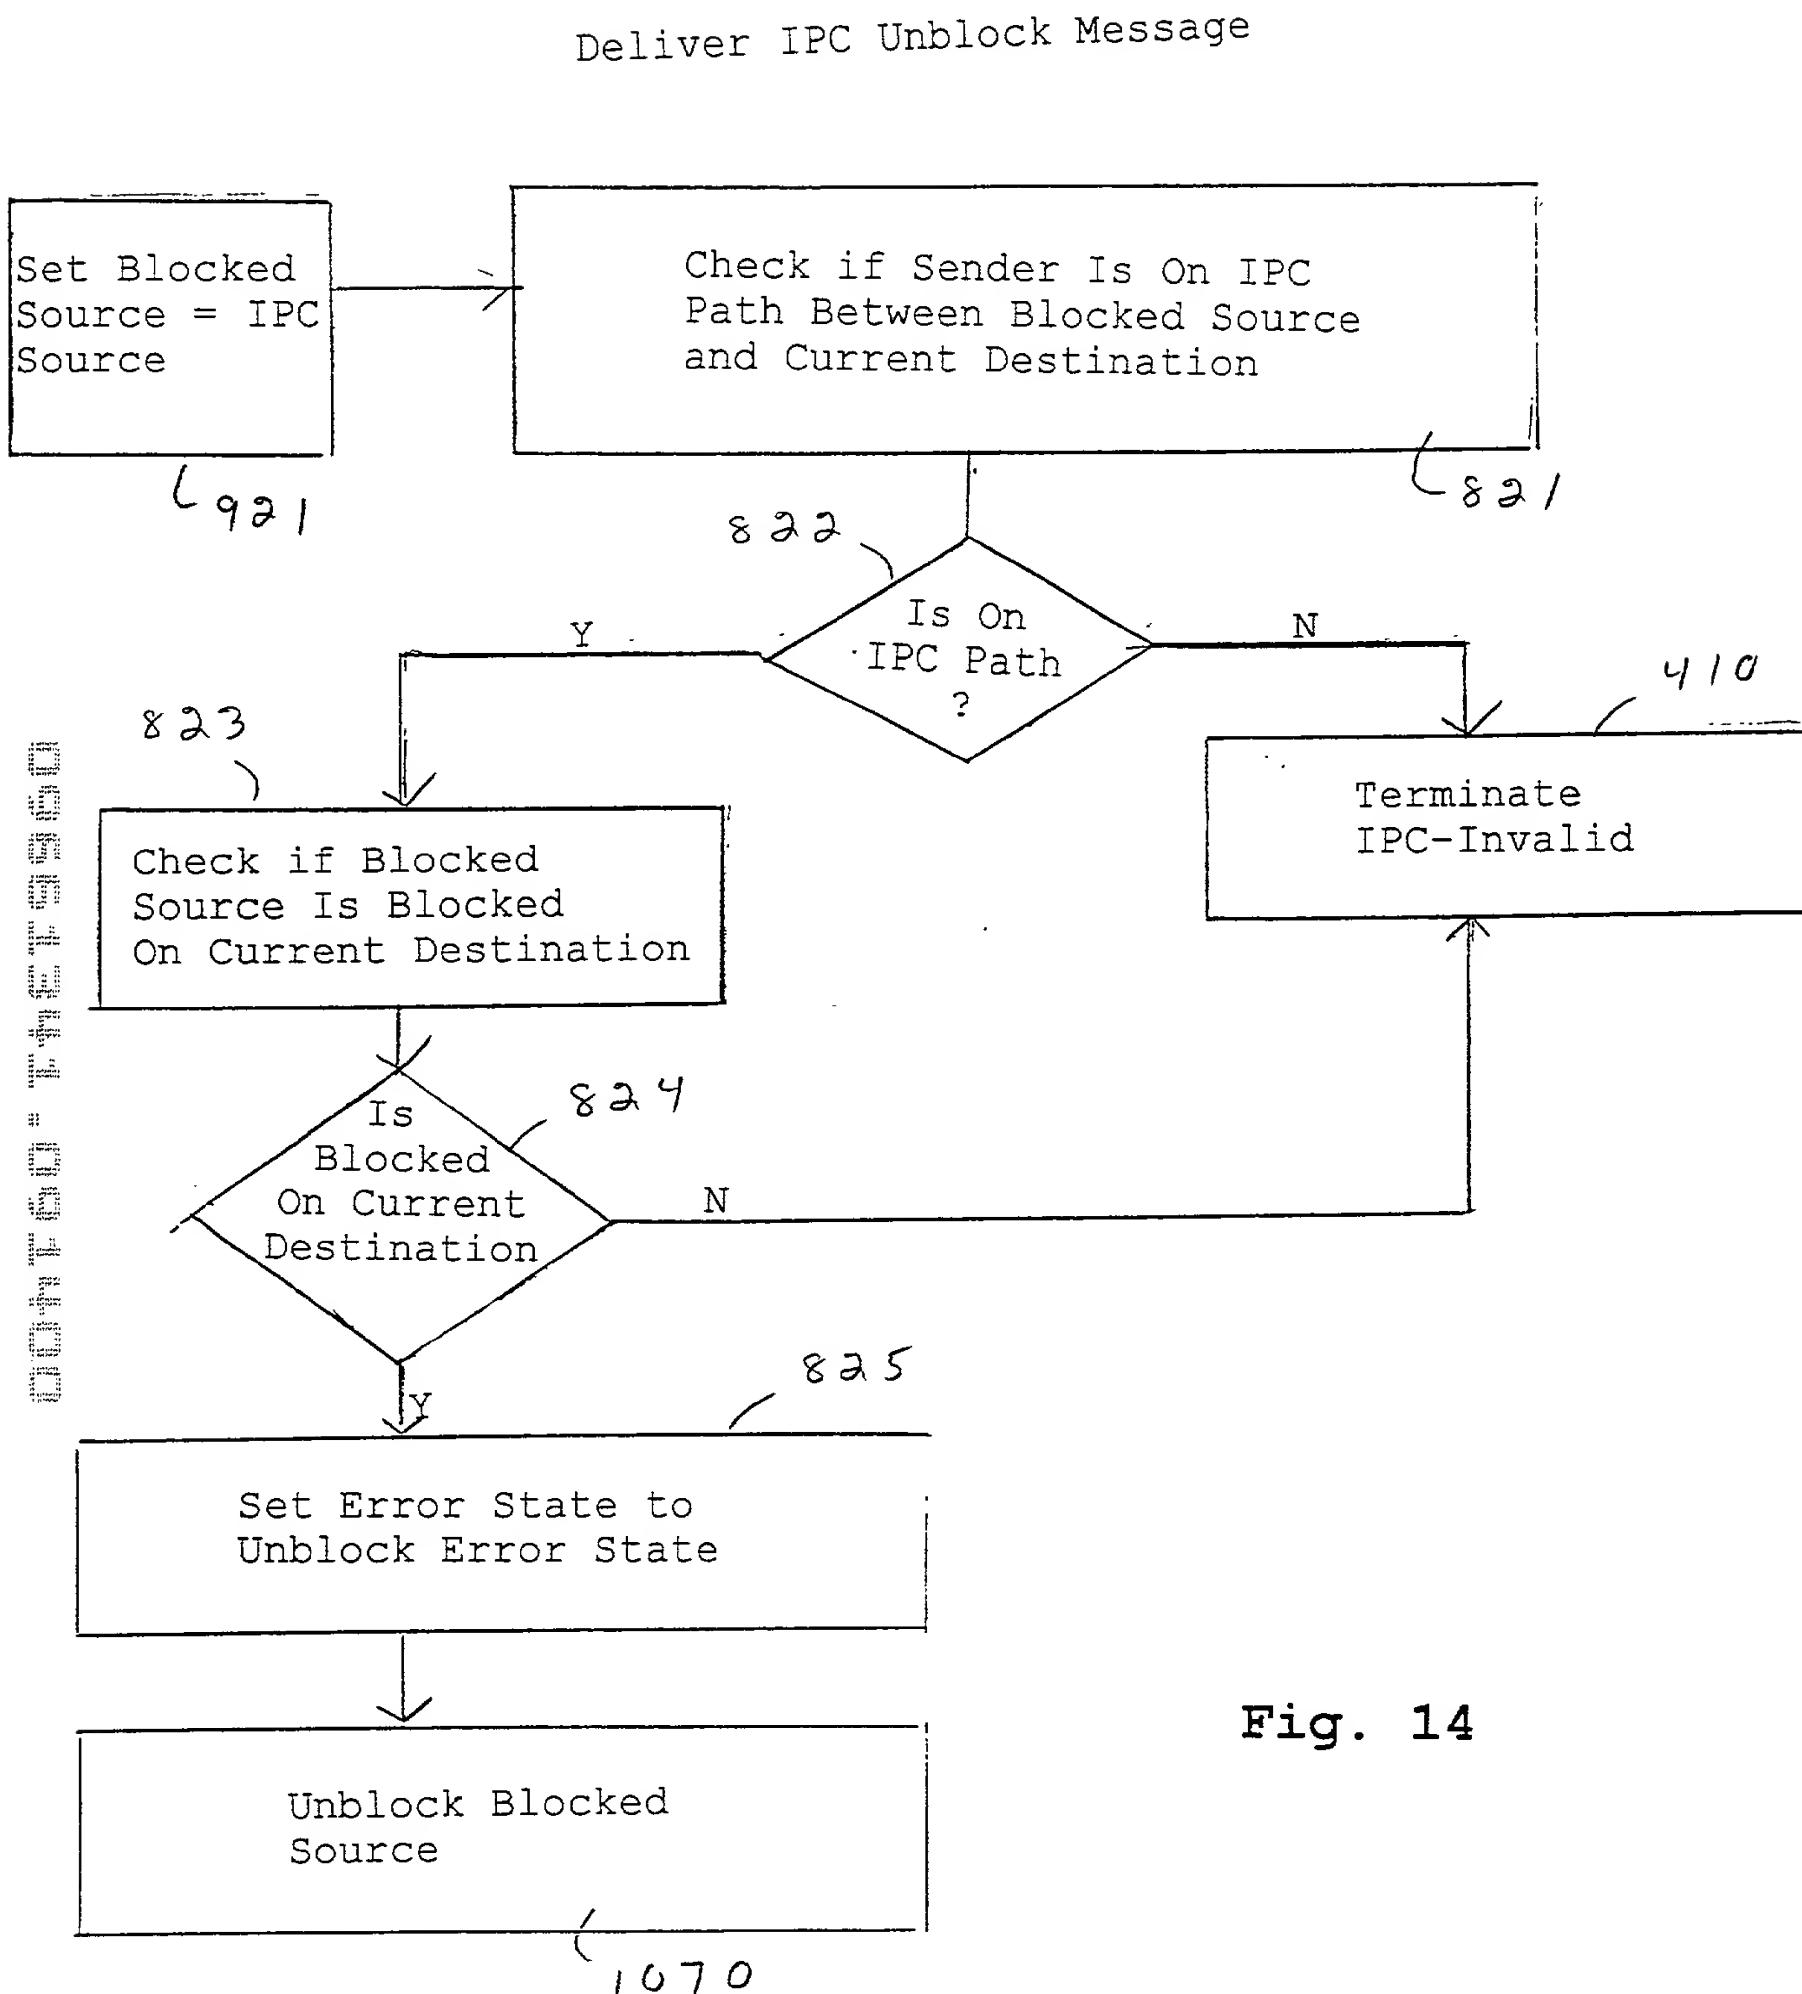


Fig. 14

Build IPC Notify Message

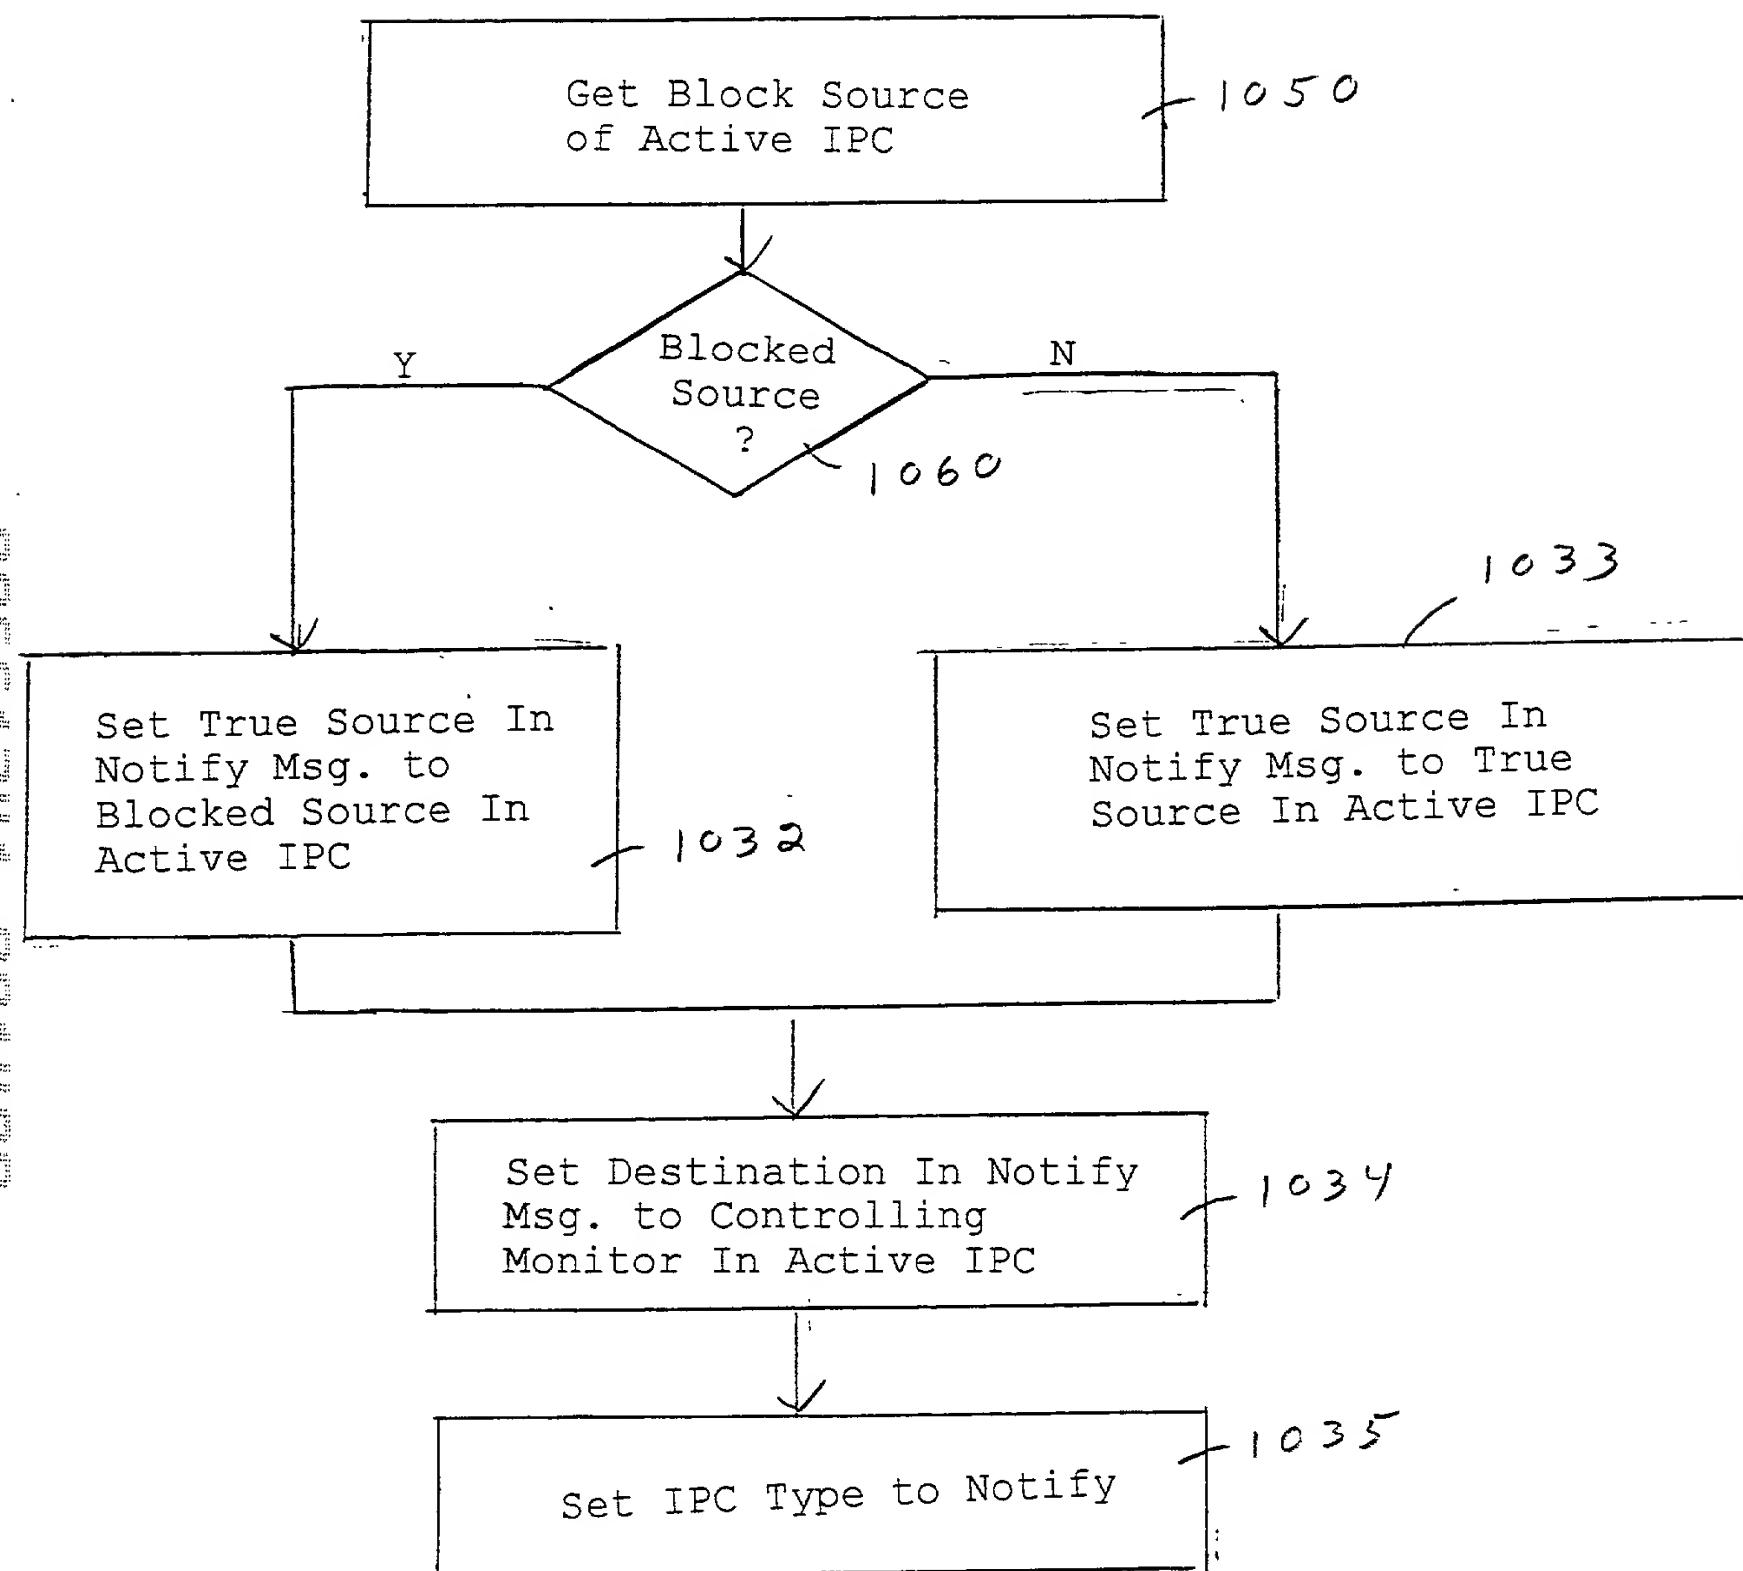


Fig. 15

Deliver Notify Message

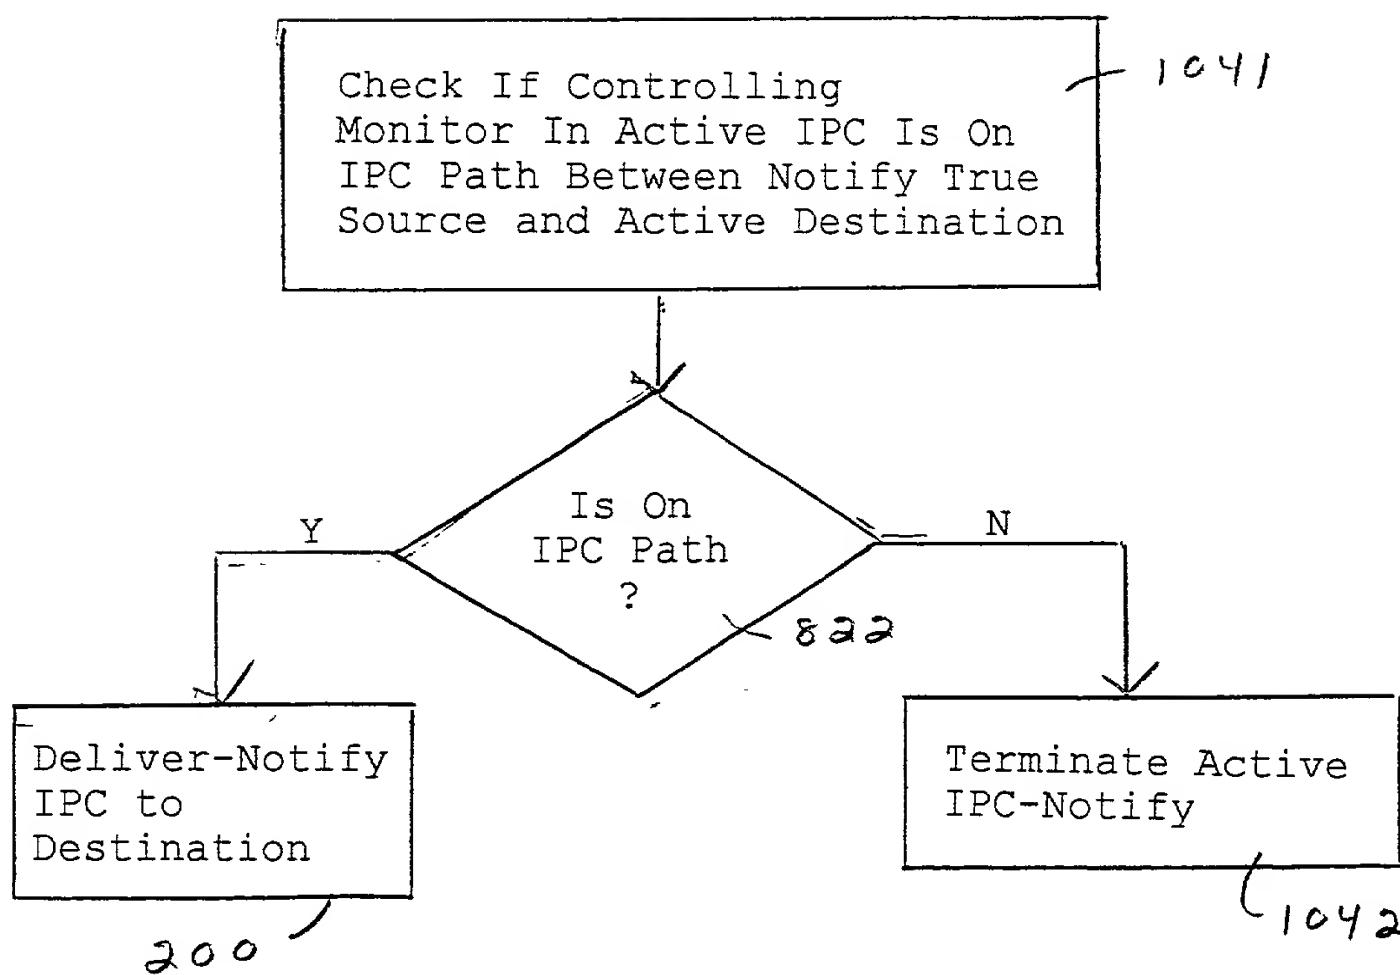


Fig. 16

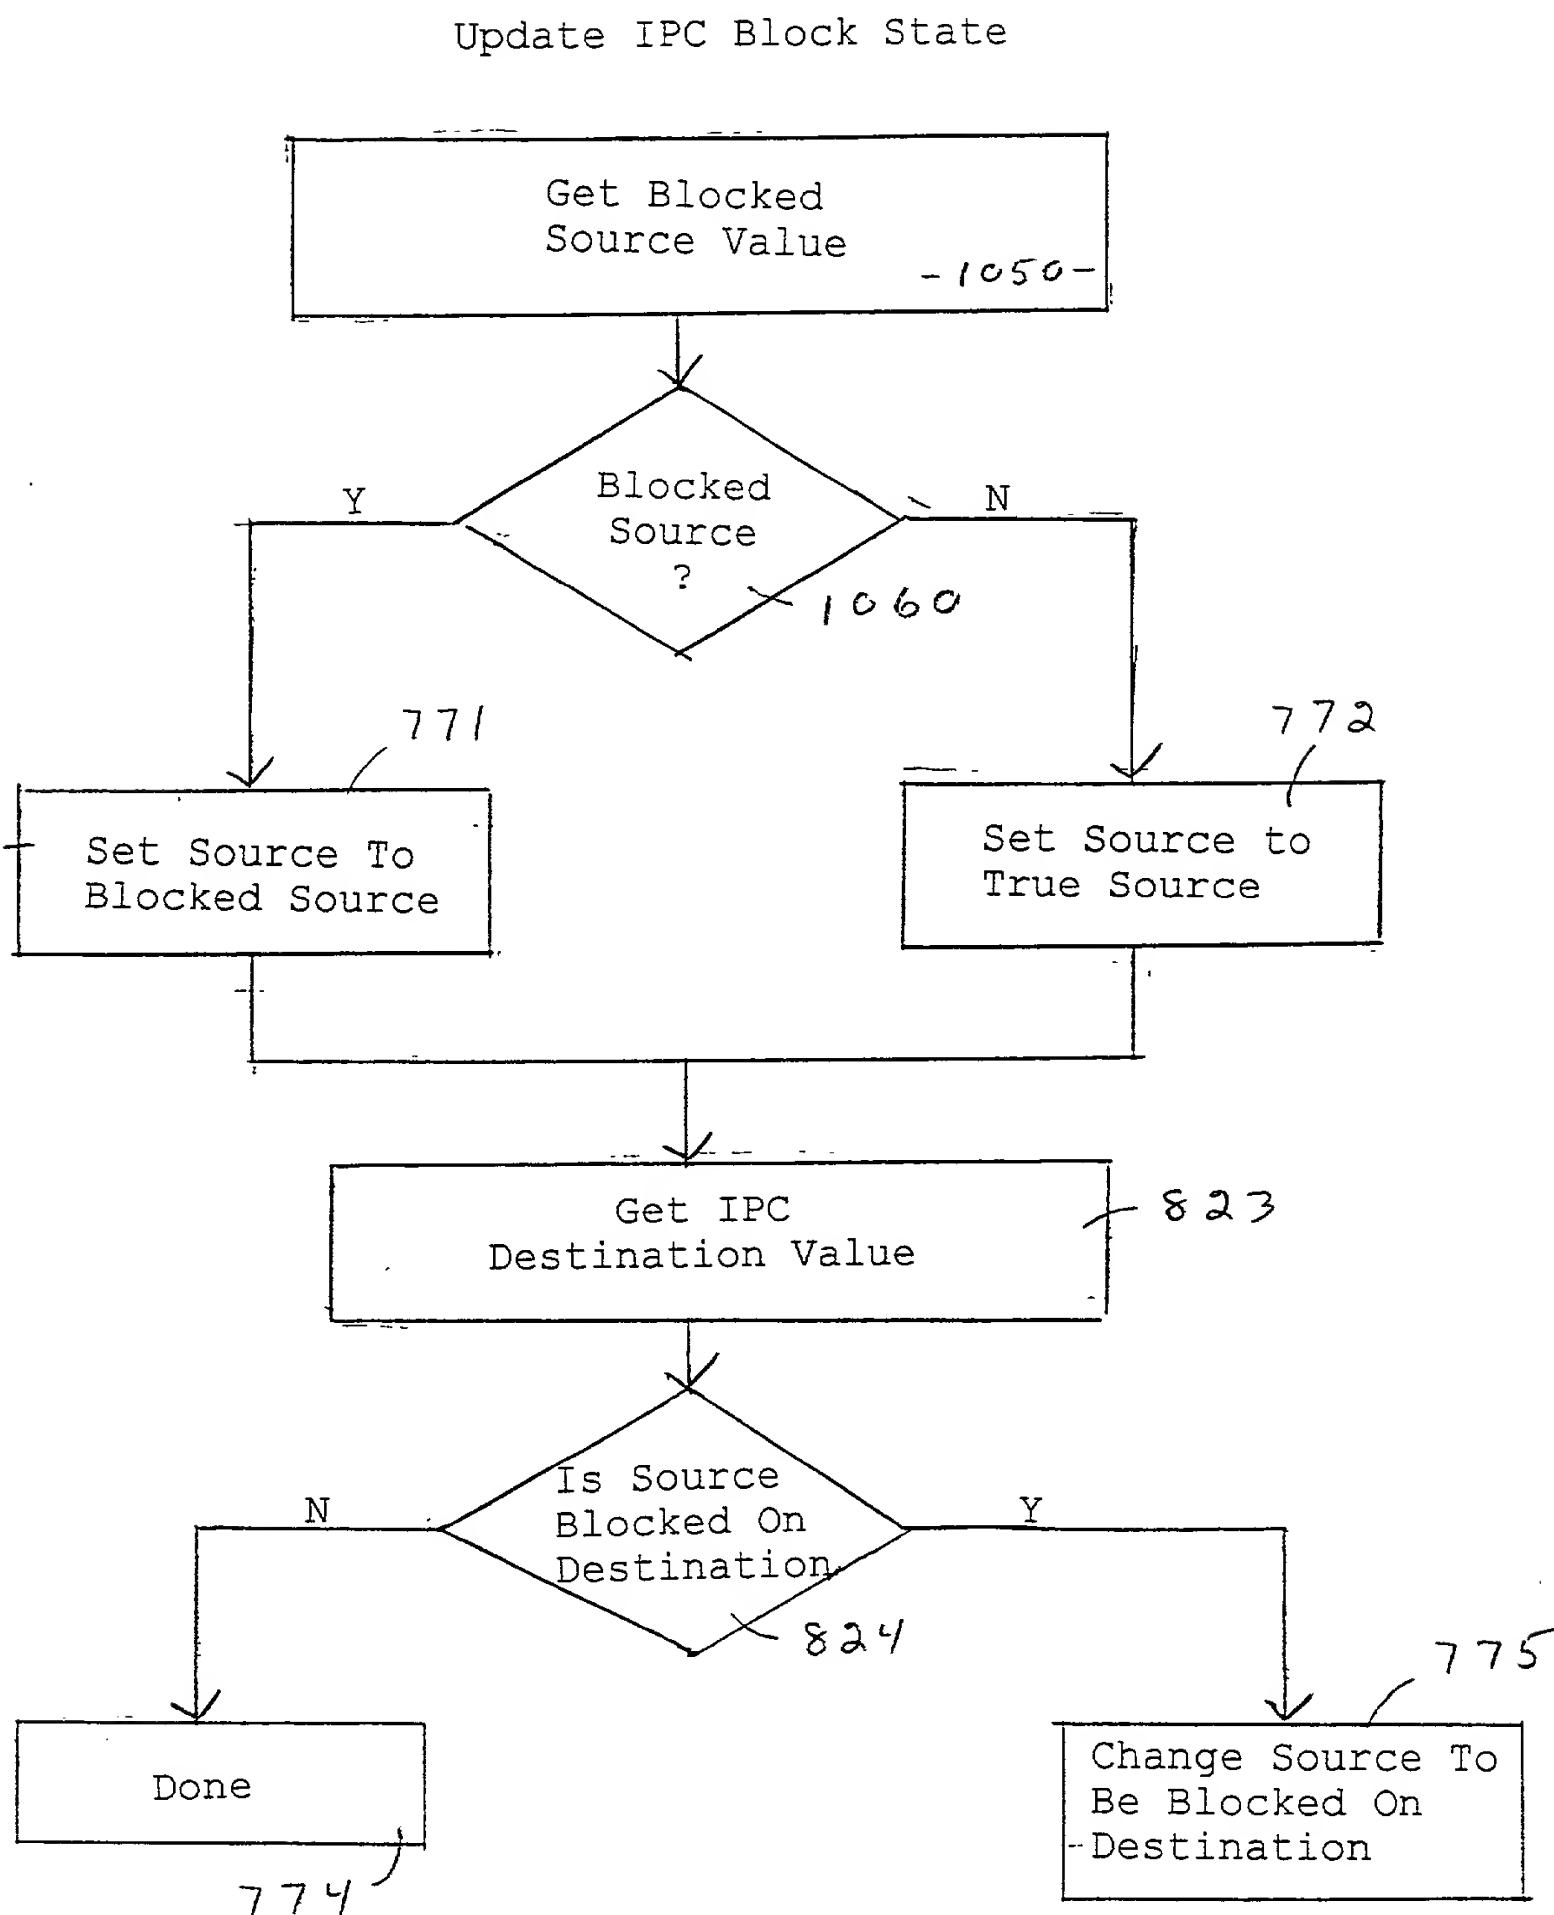


Fig. 17

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: A MECHANISM FOR SYNCHRONOUS INTERPROCESS COMMUNICATION OVER TRANSPARENT EXTERNAL MONITORS

the specification of which (check one)

is attached hereto.

was filed on _____ as United States Application Number _____

or PCT International Application Number _____

and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) or §365(b) of any foreign application(s) for patent or inventor's certificate, or §365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application, having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)	Priority Claimed
<input type="checkbox"/> (Number) <input type="checkbox"/> (Country) <input type="checkbox"/> (Day/Month/Year Filed)	<input type="checkbox"/> Yes <input type="checkbox"/> No
<input type="checkbox"/> (Number) <input type="checkbox"/> (Country) <input type="checkbox"/> (Day/Month/Year Filed)	<input type="checkbox"/> Yes <input type="checkbox"/> No
<input type="checkbox"/> (Number) <input type="checkbox"/> (Country) <input type="checkbox"/> (Day/Month/Year Filed)	<input type="checkbox"/> Yes <input type="checkbox"/> No

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

<input type="checkbox"/> (Application Number)	<input type="checkbox"/> (Filing Date)
<input type="checkbox"/> (Application Number)	<input type="checkbox"/> (Filing Date)

I hereby claim the benefit under 35 U.S.C. §120 of any United States Application(s), or §365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States, or PCT International application in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in 37 CFR §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

<input type="checkbox"/> (Application Serial No.)	<input type="checkbox"/> (Filing Date)	<input type="checkbox"/> (Status) (patented, pending, abandoned)
<input type="checkbox"/> (Application Serial No.)	<input type="checkbox"/> (Filing Date)	<input type="checkbox"/> (Status) (patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith (list name and registration number).

Manny W. Schecter (Reg. 31,722), Lauren C. Bruzzone (Reg. No. 35,802), Christopher A. Hughes (Reg. 26,914), Edward A. Pennington (Reg. 32,588), John E. Hoel (Reg. 26,279), Joseph C. Redmond, Jr. (Reg. 18,753), Douglas W. Cameron (Reg. No. 31,596), Wayne L. Ellenbogen (Reg. No. 43,602), Stephen C. Kaufman (Reg. No. 29,551), Daniel P. Morris (Reg. No. 32,053), Louis J. Percello (Reg. No. 33,206), David M. Shofi (Reg. No. 39,835), Robert M. Trepp (Reg. No. 25,933), Paul J. Otterstedt (Reg. No. 37,411) and Louis P. Herzberg (Reg. No. 41,500) and Marian Underweiser (Reg. No. 46,134).

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

Send Correspondence to: Richard L. Catania, Scully, Scott, Murphy & Presser

400 Garden City Plaza, Garden City, New York 11530

Direct Telephone Calls to: (name and telephone number) Richard L. Catania, (516) 742-4343

Trent Ray Jaeger

Full name of sole or first inventor

Inventor's Signature

Date

30 Wells Avenue, Croton-on-Hudson, New York 10520

Residence

United States of America

Citizenship

Same as residence

Post Office Address

Jonathon Earnshaw Tidswell

Full name of second joint inventor, if any

Inventor's signature

Date

Residence

Citizenship

Post Office Address